NORTEL NETWORKS

Enterprise Edge IP Telephony Configuration Manual

© 1999 Nortel Networks

PO908509 Issue 01

Contents

Chapter 1 Overview 7

About this document 8
System Functions 8
Dialing Plan Support 9
Enterprise Edge IP telephony and M1 networking 9
Toll bypass with VoIP Gateway 12
Network Quality of Service 16
Network Monitoring 17
Quality of service parameters 17
Fallback to circuit-switched voice facilities 18
Network Performance Utilities 18
Codecs 19
Silence compression 22
Echo cancellation 23
Non-linear processing 24
Jitter buffer 24
Fax calls 25
Alarm Notification 26

Chapter 2 Engineering guidelines 27

Introduction 27
Enterprise Edge IP telephony 27
Overview 28
Enterprise Edge VoIP Gateway bandwidth engineering 29
Multiple network interfaces 30
Method 1 31
Method 2 32
LAN engineering 32
Silence compression 33
WAN engineering 35
Assessing WAN link resources 35
Link utilization 36
Estimating network loading due to IP telephony traffic 36
Other intranet resource considerations 38
Setting QoS 39
Measuring Intranet QoS 40
Measuring end-to-end network delay 40
Measuring end-to-end packet loss 41
Recording routes 41
Adjusting `ping` measurements 42

Measurement procedure 43
Other measurement considerations 43
Further network analysis 44
Components of delay 44
Reducing link delay 45
Reducing hop count 46
Routing issues 47
Implementing QoS in IP networks 47
Traffic mix 48
TCP traffic behavior 49
Enterprise Edge Router QoS Support 49
Implementing the network 49
LAN engineering 49
Fallback threshold 52
Post-installation network measurements 53
Setting IP telephony QoS objectives 53
Intranet QoS monitoring 54
User feedback 54
Dialing plan 55
IP telephony and M1 networking 55
Toll bypass with IP telephony 58
Core telephony services configuration 62

Chapter 3 Engineering checklist 65

Chapter 4 Installation 67

Installation Roadmap 67
Configuring the local gateway 67
Adding a remote gateway 68

Chapter 5 Configuration 73

User Interface Overview 74
Local gateway configuration 75
Remote gateway configuration 78
Core telephony services configuration 79
Configuration of fallback to conventional circuit-switched facilities 80

Chapter 6 Maintenance 81

Quality of Service Monitor 81
Quality of Service Status 81
Using the QoS Monitor pull-down View menu 81
Operational Statistics 81
Backup and Restore Procedures 81

Chapter 7 Interoperability 83

interoperability considerations 83

Asymmetrical media channel negotiation 84

No feedback busy station 84

Glossary 85

Index 87

Overview

The Enterprise Edge VoIP Gateway reduces customers' communication costs by routing voice traffic over private Internet Protocol (IP) networks as part of the Enterprise Edge product portfolio. Enterprise Edge uses IP telephony to link multiple sites together using an existing corporate data network. The IP trunks are an integral part of the telephony services. IP telephony is transparent to users.

Enterprise Edge provides IP telephony capability. IP telephony involves the conversion of voice from its traditional telephony format (continuous analog or digital signal) into a digital packet format that can be transported over an intranet.

IP telephony operates on an installed corporate IP network. IP telephony requires a well managed intranet, rather than the internet. The private IP network facilities must have under-utilized bandwidth on the private Wide Area Network (WAN) backbone. The Engineering guidelines chapter of this guide contains information on determining if your corporate IP network can support IP telephony. A keycode controls the number of supported IP ports.

IP telephony uses a Web-based browser for configuration. See the Configuration chapter of this guide for information on how to configure IP telephony.

VoIP Gateway supports ITU-H.323v2 gateway operation. VoIP Gateway uses standard Digital Signal Processor (DSP) voice coding. See the *Enterprise Edge Programming Operations Guide* for information on DSP. VoIP Gateway supports compression algorithms (codecs) such as G.711, G.723, and G.729. See Codec types in the Engineering guidelines chapter for information on codecs.

VoIP Gateway monitors the data network and reroutes calls to the conventional circuit-switched voice facilities if Quality of Service (QoS) over the data network declines. This Fallback to Conventional Circuit-Switched Voice Facilities feature allows the system and installer to determine the acceptable QoS over the data network. The customer can configure QoS parameters according to their requirements. See the Quality of service parameters and Configuration of fallback to conventional circuit-switched facilities sections in the Configuration chapter for information on configuring the QoS parameters. If the quality falls below the expected level of QoS, the regular circuit-switched voice facilities route is selected until the QoS returns to an acceptable level.

About this document

This guide provides information on the Enterprise Edge VoIP Gateway. This guide is addressed to both telecom and datacom engineers who are going to design and implement the network. It is assumed that the telecom engineer is familiar with engineering the Enterprise Edge product portfolio, and obtaining system voice and fax traffic statistics. It is assumed that the datacom engineer is familiar with the intranet architecture, LAN implementation, tools for collecting and analyzing data network statistics, and data network management systems. The terms installer and administrator used in this document refer to the person in either the telecom or datacom engineering role. This guide contains the following sections:

- Engineering guidelines
- Engineering checklist
- Installation
- Configuration
- Maintenance
- Interoperability

System Functions

Enterprise Edge VoIP Gateway uses IP telephony to provide least cost routing of voice traffic through a corporate intranet. VoIP Gateway provides the following:

- Basic calls with answer and disconnect supervision
- Direct Inward Dial (DID) and Direct Outward Dial (DOD)
- Calling name and number
- VoIP Gateway to M1-ITG capability
- ITU-H.323 v2 compatible gateway
- Economical bandwidth use through voice compression
- Economical bandwidth use through silence compression
- Quality of Service (QoS) monitoring of gateways
- Circuit-switched voice facilities fallback capability

The core telephony service offered through Enterprise Edge treats the Enterprise Edge VoIP Gateway as a trunk. The IP trunk uses the trunking and routing functionality of the Enterprise Edge product portfolio. The IP trunks are an integral part of the Enterprise Edge product portfolio.

VoIP Gateway trunks are supervised trunks with answer and disconnect supervision. The VoIP Gateway supports voice and fax calls. See the Engineering guidelines chapter for more information about fax calls. VoIP Gateway does not support modem calls.

The IP telephony gateway allows communication with other supported gateways and H.323 v2 gateways through trunk calls. The IP telephony gateway supports Direct Routed communication. The local gateway performs the address resolution. The local gateway maintains the remote gateway table.

Dialing Plan Support

Dialing plan configuration allows the customer to set up the routing tables to route calls to appropriate destinations based on the dialed digits.

Routing codes and the destination code table allow the core telephony services on the Enterprise Edge to decide which trunking facilities are used for calls and when they are used.

Enterprise Edge has two main areas of configuration: the destination codes in the core telephony services and the destination digits in the remote gateway configuration table. The destination digits allow VoIP Gateway to route calls to the appropriate intranet destination based on the leading dialed digits. The destination code tables route calls to the appropriate trunks based on the leading dialed digits.

See the Configuration chapter for details on configuring destination digits and destination codes.

The dialing plans for all VoIP Gateways connected to the corporate intranet need to be coordinated so that calls can be made between gateways as required.

Enterprise Edge IP telephony and M1 networking

This example shows a private network composed of one central Meridian 1, and two smaller sites with Enterprise Edge systems connected over IP trunks through a corporate IP network. This could represent a large head office (with the Meridian 1) connected to several smaller branch offices.

In this network, only the head office has trunks connected to the public network. The branch offices access the public network using IP trunks to the head office. This configuration allows for cost savings by consolidating the public access trunks. Users at all three locations access the public network by dialing ‘9’, followed by the public number. For example, a user in the west end branch might dial 9-555-1212 (for a local call) or 9-1-613-555-1212 (for a long distance call). These public calls are routed to the Meridian 1 by the Enterprise Edge routing table. Routing tables at the Meridian 1 will then select an appropriate public facility for the call.

Private network calls are made by dialing a 4-digit private network DN. For example, if a user in the west end branch wishes to call a user in the east end branch within the private network, they dial 6221.

Figure 1 Enterprise Edge and M1 networking overview

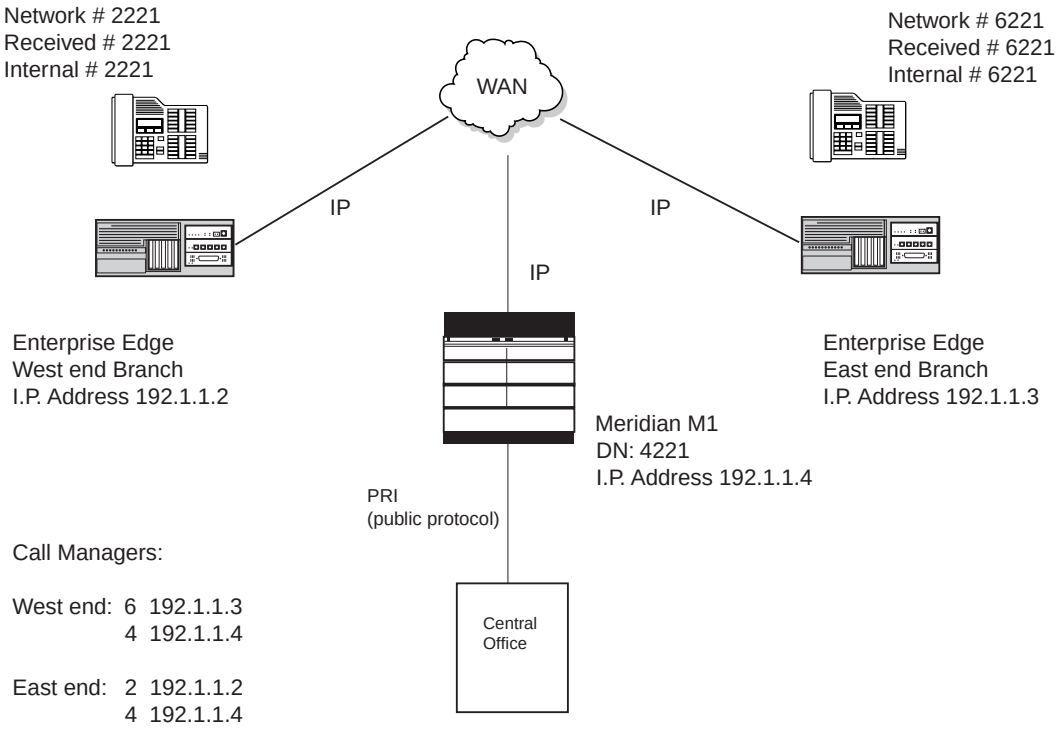

Note: The quality of the IP trunk connection is assessed during initial call setup, and if the quality is poor, Enterprise Edge will try to find an alternate route to complete the call (fallback) based on the programming definitions in the routing table. For simplicity, this example does not show programming for fallback. In this example, if the quality of the IP connection is considered too low during the call setup phase, the call would fail. For an example of fallback programming, refer to the section, “Toll bypass with VoIP Gateway” on page 12.

Note: Enterprise Edge VoIP Gateway requires a keycode. After entering the keycode for Enterprise Edge IP Telephony, perform a warm reset by following the procedure in the Maintenance section of the *Enterprise Edge Programming Operations Guide*.

In the table that follows, private network routing information is highlighted in gray. Public network routing information is shown in white.

The gateways examine the dialed digits and route the call to the corresponding IP address.

Heading	Parameter	Setting
West End office:		
Trunk/Line Data	Line 241	Target line
	Received #	2221
Line Access	Set 2221	L241:Ring only
	Line pool access	Line pool A
To Head office (M1):		

Heading	Parameter	Setting
Service/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
	Destination Code	4
	Normal route	001
	Absorb	None
To East End:		
Service/Routing Service	Destination Code	6
	Normal route	001
	Absorb	None
To Public Network:		
Service/Routing Service	Route	002
	Use	Pool A
	External #	(leave blank)
	DN type	Public
	Destination Code	9
	Normal route	002
	Absorb	None
East End office:		
Trunk/Line Data	Line 241	Target line
	Received #	6221
Line Access	Set 6221	L241:Ring only
	Line pool access	Line pool A
To Head Office: (M1)		
Service/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
	Destination Code	4
	Normal route	001
	Absorb	None
To West End:		
Service/Routing Service	Destination Code	2
	Normal route	001
	Absorb	None
To Public Network:		
Service/Routing Service	Route	002
	Use	Pool A

Heading	Parameter	Setting
	External #	(leave blank)
	DN type	Public
	Destination Code	9
	Normal route	002
	Absorb	None

In this example, outgoing public network calls dialed from a Enterprise Edge Voice Soution set are passed to the Meridian M1, and the Meridian M1 is responsible for seizing a public trunk. For this reason, the ‘9’ prefix is left in the number passed to the Meridian 1.

Note: Ensure that Line Pool A is used for IP trunks.

In order for the digit counting algorithm for outgoing IP calls to take into account this extra digit, the Private Network Access Code must be set to ’9’ on each Enterprise Edge system.

The Meridian M1 must recognize incoming 2xxx and 6xxx DID calls, and route the call over IP trunks to either the East or West end offices.

The Meridian M1 must recognize numbers starting with ‘9’ as public numbers, whether the numbers are dialed by Meridian M1 users or by Enterprise Edge Voice users.

Toll bypass with VoIP Gateway

This example shows a private network composed of one Enterprise Edge in Toronto and one Enterprise Edge in Ottawa, connected over IP trunks through a corporate IP network.

In this network, each Enterprise Edge has a PRI trunk to the Central Office, and IP trunks to the other Enterprise Edge. Calls from the Toronto system to the Ottawa system and the Ottawa public network are made over IP trunks with fallback to the PRI trunks when IP trunks are congested. This configuration allows for cost savings by using the corporate IP network whenever possible, thereby bypassing toll charges that would be incurred by using the public network.

Note: When a call gets rerouted over the PSTN due to congestion, the user may see a prompt "Expensive route." The warning indicates that toll charges may be applied to this call.

Users at both locations access the public network by dialing ’9’, followed by the public number. For example, a user in Toronto might dial 9-555-1212 (for a local call), or 9-1-613-555-1212 (for a long distance call to Ottawa). Local calls would be sent directly to the Central Office over PRI trunks. Long distance calls to Ottawa would be sent over IP trunks; the Ottawa system would tandem these calls to the local Central Office over PRI trunks.

Private network calls are made by dialing a 4-digit private network DN. For example, if a user in Toronto wants to call a user in Ottawa within the private network, they dial 6221.

Note: Enterprise Edge VoIP Gateway requires a keycode. After entering the keycode for Enterprise Edge IP Telephony, perform a warm reset by following the procedure in the Maintenance section of the *Enterprise Edge Programming Operations Guide*.

Figure 2 Toll bypass overview

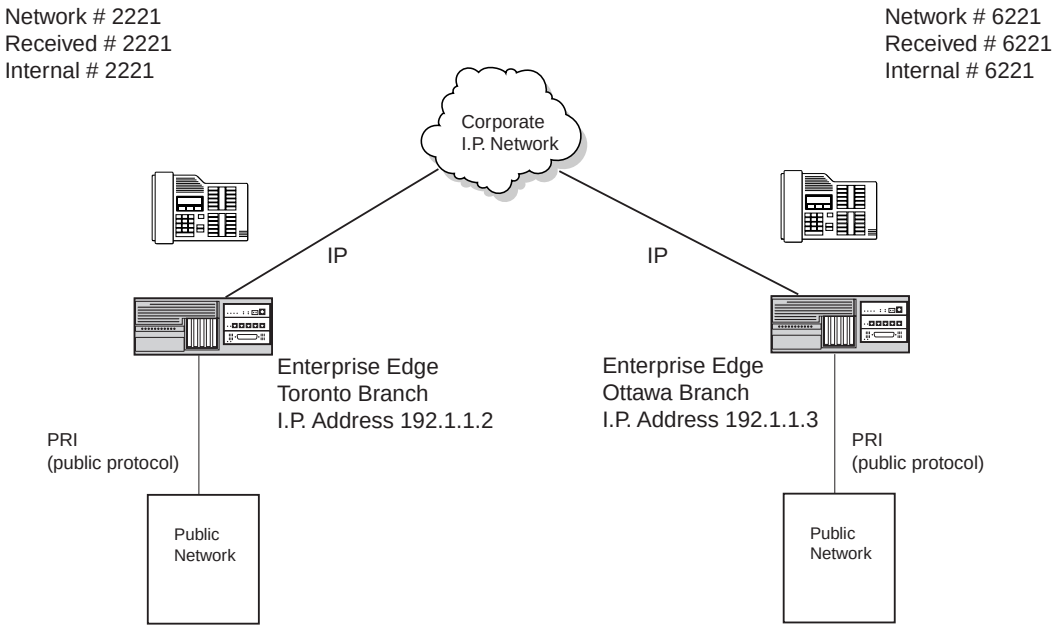

The Gateway at the Toronto office examines the dialed digits and determines that it should be routed to the IP address corresponding to the Ottawa office. The Ottawa office receives the call, sees that the leading digit(s) match its Private Network Access Code, and uses a destination code to route the call over its public trunks to the PSTN.

This is a simplified example where only calls to the 613 Area Code are routed by the Ottawa node. In a real world configuration, it would also be desirable to handle Area Codes that are ‘close’, for example Montreal: 514.

In the table that follows, private network routing information is highlighted in gray. Public network routing information is shown in white.

Heading	Parameter	Setting
Toronto office:		
Lines/Trunk/Line Data	Line 241	Target line
	Received #	2221
Terminals & sets/Line Access	Set 221	L241:Ring only
	Line pool access	Line pool A
		Line pool PRI-A

Heading	Parameter	Setting
Calls to Ottawa office:		
Services/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
Services/Routing Service	Route	002
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Private
Services/Routing Service	Destination Code	6
	Schedule 4	001
	Absorb	None
	Normal route	002
	Absorb	None
Calls to Ottawa Public Network:		
Services/Routing Service	Route	003
	Use	Pool A
	External #	(leave blank)
	DN type	Public
	Route	004
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Public
	Destination Code	91613
	Normal route	004
	Absorb	1
	Schedule 4	003
	Absorb	None
To Public Network:		
Services/Routing Service	Destination Code	9161A
	Normal route	004
	Absorb	1
	Destination Code	916A
	Normal route	004
	Absorb	1
	Destination Code	91A
	Normal route	004
	Absorb	1
	Destination Code	9A
	Normal route	004
	Absorb	1

Heading	Parameter	Setting
Ottawa office:		
Trunk/Line Data	Line 241	Target line
	Received #	6221
Line Access	Set 6221	L241:Ring only
	Line pool access	Line pool A Line pool PRI-A
To Toronto office:		
Services/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
	Route	002
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Private
	Destination Code	2
	Normal route	002
	Absorb	None
	Schedule 4	001
	Absorb	None
To Toronto Public Network:		
Services/Routing Service	Route	003
	Use	Pool A
	External #	(leave blank)
	DN type	Public
Services/Routing Service	Route	004
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Public
	Destination Code	91416
	Normal route	004
	Absorb	1
	Schedule 4	003
	Absorb	None
To Public Network:		
Services/Routing Service	Destination Code	9141A
	Normal route	004
	Absorb	1
	Destination Code	914A
	Normal route	004
	Absorb	1

Heading	Parameter	Setting
	Destination Code	91A
	Normal route	004
	Absorb	1
	Destination Code	9A
	Normal route	004
	Absorb	1

The implications on the configuration on each node are:

- each node must have the Private Network Access Code set to the value 9.
- each node must have destination code(s) that match the Private Network Access Code plus digits corresponding to calls terminating in the local PSTN. For example, if the Private Network Access Code is ‘9’, the node in Ottawa would require a destination code of ‘91613’. Similarly, Toronto would require the following destination code: 91416.

  Note: Ensure that Line Pool A is used for IP trunks.
- To allow for fallback to PRI trunks when the IP trunks are congested, you must also program the following Routing service settings:
- Set the start and end times for Sched 4 to 1:00 so that IP calls can be made 24 hours a day.
- Program the Sched 4 Service setting to Auto and enable overflow routing by changing the Overflow setting to Y (Yes).
- A control set must be defined for all sets on the system that make calls over IP trunks. See the *Enterprise Edge Programming Operations Guide* for more information.

You must program Remote Packages so that the IP trunks in Pool A can access the lines in Pool PRI-A in a toll bypass scenario. In other words, you must give package 01 access to pool PRI-A and you must assign package 01 to all IP trunks. For more information, see the *Enterprise Edge Programming Operations Guide*.

Network Quality of Service

Enterprise Edge VoIP Gateway uses a method similar to ITU-T Recommendation G.107, the E-Model, to determine the voice quality. This model evaluates the end-to-end network transmission performance and outputs a scalar rating “R” for the network transmission quality. The packet loss and latency of the end-to-end network determine “R”. The model further correlates the network objective measure “R”, with the subjective QoS metric for voice quality, MOS or the Mean Opinion Score.

This model serves as an effective traffic shaping mechanism by invoking the *Fallback to Circuit-Switched Voice Facilities* feature at call set up to avoid quality of service degradation. New calls fall back when the configurable MOS values for all codecs fall below the threshold.

The model accounts for compression characteristics of the codecs. Each codec delivers a different MOS for the same network quality.

Network Monitoring

The VoIP Gateway network monitoring function measures the quality of service between the local and all remote gateways on a continuous basis. The network monitoring function exchanges UDP probe packets between all monitored gateways to collect the network statistics for each remote location. All the packets make a round trip from the Sender to Receiver and back to the Sender. From this information, the latency and loss in the network for a particular location are calculated.

It may take about 3 mins before the VoIP Gateway monitoring function reacts to marginal changes in the network condition. Fallback can be due to any of the following reasons:

- Bad network conditions.
- The remote gateway is out of service.
- No network connection.

Note 1: Quality of Service monitoring is not supported for non-Enterprise Edge product locations and must be disabled.

Note 2: The Quality of Service threshold is configurable per remote gateway.

Note 3: Fallback is triggered for all new originating calls if the QoS of any monitored gateway falls below its threshold.

Note 4: The fallback decision is made only at the originating gateway using the QoS thresholds monitored at the originating gateway for the destination gateway.

VoIP Gateway allows for manual configuration of QoS thresholds depending on the customer trade-off between cost and voice quality. The Engineering guidelines chapter provides the necessary guidelines to effectively weigh the trade-off and determine the quality of service that can be supported for any given network.

Quality of service parameters

Quality of Service is largely dependent on end-to-end network performance and available bandwidth. A number of parameters determine the VoIP Gateway QoS over the data network.

Packet loss

Packet loss is the percentage of packets that do not arrive at their destination. Packet loss is caused by transmission equipment problems, and high delay and congestion. In a voice conversation, packet loss is heard as gaps in the conversation. Some packet loss, less than 5%, may be acceptable without too much degradation in voice quality. Sporadic loss of small packets may be more acceptable than infrequent loss of large packets.

Packet delay

Packet delay is the time between when a packet is sent and when it is received. The total packet delay time consists of fixed and variable delay. Variable delay is the more manageable delay, since fixed delay is dependent on the network technology itself. Variable delay is caused by the particular network routing of packets. The gateway should be as close as possible to the network backbone (WAN) with a minimum number of hops, to minimize packet delay and maximize voice quality.

Delay variation (jitter)

The amount of variation in packet delay is referred to as delay variations, or jitter. Jitter affects the ability of the receiving gateway to assemble voice packets received at irregular intervals into a continuous voice stream.

Fallback to circuit-switched voice facilities

If the measured Mean Opinion Score (MOS) for all codecs falls below the configured threshold for any monitored gateway, the Fallback to Conventional Circuit-switched services is triggered. This feature reroutes calls to alternate trunks such as the Public Switched Telephone Network (PSTN). The feature reroutes calls until the network QoS improves. When the QoS meets or exceeds the threshold, calls route over the IP network.

The fallback feature can be disabled in the Local Gateway Configuration. If the fallback feature is disabled, calls are sent over the IP telephony trunks regardless of the QoS. The fallback feature is only in effect at call setup. A call in progress will not fall back if the QoS degrades.

Network Performance Utilities

Two common network utilities, `Ping` and `Traceroute`, are described below. These utilities provide a method to measure quality of service parameters. Other utilities can be used to find more information about VoIP Gateway network performance.

Note 1: Since data network conditions can vary at different times, collect performance data over at least a 24 hour time period.

Note 2: Performance utilities should be used to measure performance from each gateway to every other gateway.

Ping

`Ping` (Packet InterNet Groper) sends an ICMP (Internet Control Message Protocol) echo request message to a host, expecting an ICMP echo reply to be returned. This allows the round trip time to a particular host to be measured. By sending repeated ICMP echo request messages, percent packet loss for a route can also be measured.

Traceroute

`Traceroute` uses the IP TTL (time-to-live) field to determine router hops to a specific IP address. A router must not forward an IP packet with a TTL field of 0 or 1. It must instead throw away the packet and return to the originating IP address an ICMP “time exceeded” message.

`Traceroute` uses this mechanism by sending an IP datagram with a TTL of 1 to the specified destination host. The first router to handle the datagram will send back a “time exceeded” message. This identifies the first router on the route. The `traceroute` sends out a datagram with a TTL of 2.

This will cause the second router on the route to return a “time exceeded” message and so on until all hops have been identified. The `traceroute` IP datagram will have a UDP Port number unlikely to be in use at the destination (usually > 30,000). This will cause the destination to return a “port unreachable” ICMP packet. This identifies the destination host.

`Traceroute` can be used to measure round trip times to all hops along a route, thereby identifying bottlenecks in the network.

Codecs

The term codec refers to the voice coding and compression algorithm used by the DSP on the telephony services and the MSPECs. See the *Enterprise Edge Programming Operations Guide* for additional information on DSP and MSPEC resources.

The codec type used on a per VoIP Gateway call basis is determined at call setup. The originating gateway will indicate to the remote gateway which codec types it supports, starting with the preferred order of usage. The remote gateway, depending on its capabilities, chooses one of the codec types and continues with the call. If both ends cannot agree on a codec type, the call fails.

Therefore, it is important that all gateways in the intranet use the same codec types.

Each gateway needs to be configured with which possible codecs are available for negotiation, as well as the preferred order of usage. Given that the trade-off is quality versus bandwidth, the codecs configuration should reflect available bandwidth on the network.

The supported codec types are configured in the Local gateway configuration section. The G.711 codec provides the best audio quality but uses the greatest amount of bandwidth. The G.729 and G.723.1 codecs use less bandwidth, but reduce audio quality. The installer or administrator determines the best choice for the user and the available bandwidth on the intranet. For example, if the WAN link cannot support multiple 64 kbit/s calls, G.711 should not be configured as a supported codec.

Enterprise Edge Solutions recommends the following order for codec selection:

- G.729
- G.723.1 (6.3 kbit/s or 5.3 kbit/s)
- G.711

The G.729 codec provides the best balance of quality audio plus bandwidth savings.

Enterprise Edge VoIP Gateway supports the following codecs:

G.711

This codec delivers “toll quality” audio at 64 kbit/s. This codec is optimal for speech since it has the smallest delay, and is very resilient to channel errors. However, it consumes the largest bandwidth. North America uses G.711 μ-LAW and international markets use G.711 A-LAW.

G.729

The G.729 codec is the default and preferred codec for IP telephony. It provides near toll quality with a low delay. This codec uses compression to 8 kbit/s. Enterprise Edge VoIP Gateway supports G.729 with silence compression, per Annex B.

G.723.1

The G.723.1 codec uses the smallest amount of bandwidth. This codec uses the greatest compression, 5.3 kbit/s or 6.3 kbit/s.

The G.723.1 codec uses a different compression method than the G.729 codec. The G.723.1 method uses more DSP resources. Each MSPEC supports only one G.723.1 call. A G.711 call can run in the same MSPEC as a G.723.1 call. See the *Enterprise Edge Programming Operations Guide* for additional information.

If the G.723.1 codec is the only possible codec for a call, a trunk may not be available for the call if there are insufficient DSP resources available. All VoIP Gateway facilities will appear to be in use, even though there are DSP resources available for calls using other codec types.

Since most gateways support the G.711 codec, configure G.711 as a supported codec. The G.711 codec does not compress audio or fax. The G.711 codec supports two IP trunks on each MSPEC. See the *Enterprise Edge Programming Operations Guide* for additional information.

Silence compression

Silence compression is supported on G.723.1 and G.729, Annex B.

A key to VoIP Gateway's success in business applications is minimizing WAN bandwidth consumption. Beyond speech compression, the best bandwidth reducing technology is silence compression, also known as silence suppression. Silence compression technology recognizes the periods of silence in a conversation , and stops sending IP speech packets during those periods. Telco studies show that in a typical phone conversation, only about 36-40% of a full-duplex conversation is active. When one person talks, the other listens (this is called half-duplex). And there are significant periods of silence during speaker pauses between words and phrases.

By applying silence compression, full duplex bandwidth consumption is reduced by the same amount, freeing up bandwidth for other voice/fax or data communications. The following figure illustrates how silence compression allows two conversations to fit in the bandwidth otherwise used by one. This 50% bandwidth reduction develops over a 20-30 second period as the conversation switches from one direction to another.

To provide a more natural sound, comfort noise is added at the destination gateway during the silent periods to calls where silence compression is active. In some cases, silence compression may cause a perceived degradation in audio quality. Silence compression can be disabled. Disabling silence compression will increase bandwidth consumption.

If VoIP Gateway serves as a tandem switch in a network where some circuit-switched trunk facilities have an excessively low audio level, silence compression, if enabled, will degrade the quality of service by causing choppiness of speech. Under these conditions, silence compression should be disabled.

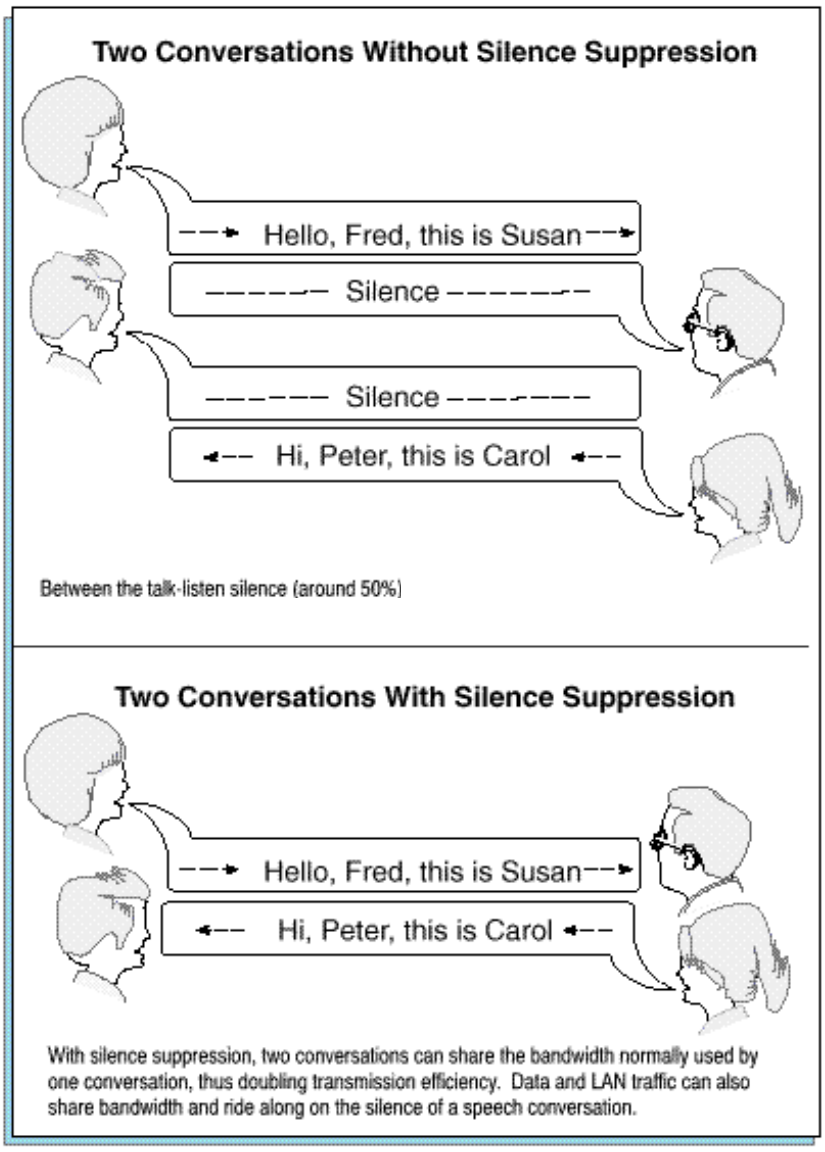
Two Conversations Without Silence Suppression
Hello, Fred, this is Susan
Silence
Silence
Hi, Peter, this is Carol
Between the talk-listen silence (around 50%)
Two Conversations With Silence Suppression
Hello, Fred, this is Susan
Hi, Peter, this is Carol
With silence suppression, two conversations can share the bandwidth normally used by one conversation, thus doubling transmission efficiency. Data and LAN traffic can also share bandwidth and ride along on the silence of a speech conversation.

Echo cancellation

When a two-wire telephone cable connects to a four-wire PBX interface or a telco central office (CO) interface, a special electrical circuit called a hybrid is used to convert between two wires and four wires. Although hybrid circuits are very efficient in their conversion ability, a small percentage of telephony energy is not converted but instead is reflected back to the caller. This is called echo.

If the caller is near the PBX or CO switch, the echo comes back so quickly it cannot be discerned. However, if the delay is more than about 10 ms, the caller can hear an echo. To prevent this, gateway vendors include special code in the DSPs that listens for the echo signal and subtracts it from the listener's audio signal. Echo cancellation is especially important for gateway vendors because the IP network delay can easily be 40–50 ms, so the echo from the far-end hybrid would be quite pronounced at the near end. Far-end echo cancellation eliminates this.

Echo cancellation sometimes causes choppiness in conversation in a low audio conversation. Although echo cancellation can be disabled, it is not recommended.

Non-linear processing

Non-linear processing (NLP) is part of echo cancellation. It improves echo cancellation by further reducing residual echo. NLP mutes background noise during periods of far-end silence and prevents comfort noise from being generated. Some listeners find muted background noise annoying. NLP can be disabled to prevent this, but with the trade-off of increased perceived echo.

Jitter buffer

A major contributor to reduced voice quality is IP network packet delay and network jitter. Network delay describes the average length of time for a packet to traverse a network. Network jitter describes the variability in arrival time of a packet. Delay is like the average, jitter is like the standard deviation. Both are important in determining voice quality.

To allow for variable packet arrival time and still produce a steady out-going stream of speech, the far-end gateway does not play out the speech as soon as the first packet arrives. Instead, it holds it for a certain time in part of its memory called the jitter buffer, and then plays it out. The amount of this hold time is the measure of the jitter buffer, e.g., a 50 ms hold time implies a 50 ms jitter buffer.

As the network delay (total time, including codec processing time) exceeds about 200 ms, the two speakers will increasingly adopt a half-duplex communications mode, where one speaks, the other listens and pauses to make sure the speaker is done. If the pauses are ill timed, they end up "stepping" on each other's speech. This is the problem that occurs when two people converse over a satellite telephony connection. The result is a reduction in perceived voice quality.

When a voice packet is inordinately delayed and does not arrive at the far-end in time to fit into the voice stream going out of the far-end gateway, it is discarded, and the previous packet is replayed. If this happens too often, or twice in a row, the listener will perceive reduced voice quality.

The jitter buffer hold time adds to the overall delay, so if the network has high jitter, the overall effect will be a long perceived delay in the voice stream. For example, a network might have a moderately average delay of 50 ms and a variability of 5 ms. The network is said to have 5 ms of jitter, a low figure. The jitter buffer hold time is only 5 ms, so the effective network total delay will only be 55 ms, still moderate.

On the other hand, if a network has a low average delay of 15 ms, but 10% of the time the delay goes out to a long 100 ms, while 90% of the time the delay is a brief 4 ms, the jitter buffer would have to be 100 ms and the total effective network delay would be 115 ms, a long delay. Network jitter can be more important than average delay in many VoIP Gateway applications.

VoIP Gateway voice calls use an adaptive jitter buffer that changes the hold time over the duration of the call. The installer or administrator configures the maximum hold time.

VoIP Gateway fax calls use a fixed jitter buffer that does not change the hold time over the duration of the call. Fax calls are more sensitive to packet loss. In situations of high jitter, increased delay (through the use of a deeper jitter buffer) is preferred. To accommodate this, VoIP Gateway provides a separate jitter buffer setting for fax calls.

Fax calls

The Enterprise Edge gateways support T.30 Group 3 fax calls. Fax calls automatically use the G.711 codec and require the associated bandwidth.

As the gateway does not know in advance that a call will carry a fax transmission, it first establishes a voice channel. The voice channel may use G.729 or G.723.1 audio compression. Upon detecting the answering fax machine's CED tone, the terminating gateway performs the following operations:

- Initiates the procedure to revert the speech path to a G.711, 64 Kbit/s clear channel.
- Disables the adaptive jitter buffer feature.
- Sets the hold time for the jitter buffer to the value specified in the Local Gateway settings to improve late IP packet tolerance.

The answering fax machine must produce its CED tone within 15 s of connection. The terminating gateway turns off CED tone detection after 15 s to prevent false tone detection during a voice call.

This method imposes the following restrictions:

- Interoperability with other IP gateways. A terminating gateway must support CED fax tone detection, and initiate the procedure as described in previous paragraphs. An originating gateway must support the H.323 Request Mode procedure, but does not need to detect fax tones. The originating gateway must additionally be capable of supporting the large G.711 packest used for fax transmission.

- In order for the gateways to revert to a G.711 clear channel, the terminating fax machine must issue a CED tone upon answering the call. Manually initiated fax transmissions, where the user at the terminating end first talks with the originating user before setting the terminating fax to receive the document, are not supported.
- Fax machines tolerate a maximum round trip delay of 1200 ms. Media processing in the the two gateways introduces a round trip delay of approximately 300 ms, in addition to the delay caused by the jitter buffer. If a 250 ms jitter buffer is used, IP latency should never exceed (1200 -(300+(2*250))) = 400 ms round trip delay, or approximately 200 ms one way.

Alarm Notification

Enterprise Edge uses the Unified Manager to capture information about its operational status.

See the Maintenance chapter for additional information.

Engineering guidelines

The engineering guidelines address the design of an IP trunk network for Enterprise Edge VoIP Gateway.The network contains the following:

- Enterprise Edge VoIP gateways
- Gateways attached to LANs
- Corporate intranet connecting the LANs

The guidelines assume that an installed corporate intranet connects the sites of the IP gateways.

Introduction

IP telephony compresses PCM voice and routes the packetized data over a private internet, or intranet, to provide virtual analog TIE trunks between gateways. Communications costs may be reduced as voice traffic is routed at low marginal cost over existing private IP network facilities with available under-utilized bandwidth on the private Wide Area Network (WAN) backbone.

This document provides guiding principles for properly designing a network of IP gateways over the corporate intranet, describe how to qualify the corporate intranet to support an IP network, and decide what required changes are needed in order to preserve the quality of voice services as much as possible when migrating those services from the PSTN. It addresses requirements for the successful integration with the customer's existing local area network (LAN). By adhering to these guidelines the designer should be able to engineer the IP so that the cost and quality trade-off is at best imperceptible, and at worst within a calculated tolerance.

Enterprise Edge IP telephony

Enterprise Edge IP telephony is designed to work on an adequately provisioned, stable LAN. Delay, delay variation or jitter, and packet loss must be minimized end-to-end across the LAN and WAN. The installer must carefully determine the design and configuration of the LAN and WAN that link the IP telephony system. If the intranet becomes congested, new calls to the IP telephony will fall back to traditional circuit-switched voice facilities so that the quality of service is not degraded for new calls.

IP telephony operates on an installed corporate IP network. IP telephony operates on a well managed intranet, rather than the internet.

Overview

Traditional networks rely on voice services such as LEC and IXC private lines. With Enterprise Edge IP telephony technology, IP telephony can now choose a new kind of delivery mechanism, one that uses packet switching over a data network, specifically a corporate intranet. The role of the IP gateway in this regard is essentially to convert steady-stream digital voice into fixed length IP packets.

New considerations come into play now when the same corporate network is expected to deliver voice traffic. The intranet introduces impairments, primarily delay, delay variation, and error, at levels that are higher than those delivered by voice networks. Delay between talker and listener changes the dynamics and reduces the efficiency of conversations, whereas delay variation and packet errors introduce glitches in conversation. Simply connecting the IP gateways to the corporate intranet without preliminary assessments may result in unacceptable degradation in the voice service; instead proper design procedures and principles must be considered.

A good design of the network must begin with an understanding of traffic, and the underlying network that will carry the traffic. There are three preliminary steps that the installer must undertake:

- Determine bandwidth requirements. The installer must estimate the amount of traffic that the Enterprise Edge product will route via the IP gateway. This in turn will place a traffic load on the corporate intranet. This is described in the Enterprise Edge VoIP Gateway bandwidth engineering section.
- Assess WAN link resources. If resources in the corporate intranet are insufficient to adequately support voice services, it is usually due to insufficient WAN resources. Assessing WAN resources is described in the Assessing WAN link resources section.
- Measuring the existing intranet’s QoS. The installer must estimate the quality of voice service the corporate intranet can deliver. The Measuring Intranet QoS section describes how to measure the prevailing delay and error characteristics of an intranet.

After the assessment phase, the installer can design and implement the IP telephony network. This design not only involves the IP telephony IP elements may also involve making design changes to the intranet.

- The Further network analysis sections provide guidelines for modifying the intranet.
- The Implementing the network section provides guidelines for integrating the IP gateway into the corporate LAN.

The following flowchart shows the design and planning decisions that should take place. Each action and decision point is addressed in this document.

Figure 3 IP Telephony network engineering process

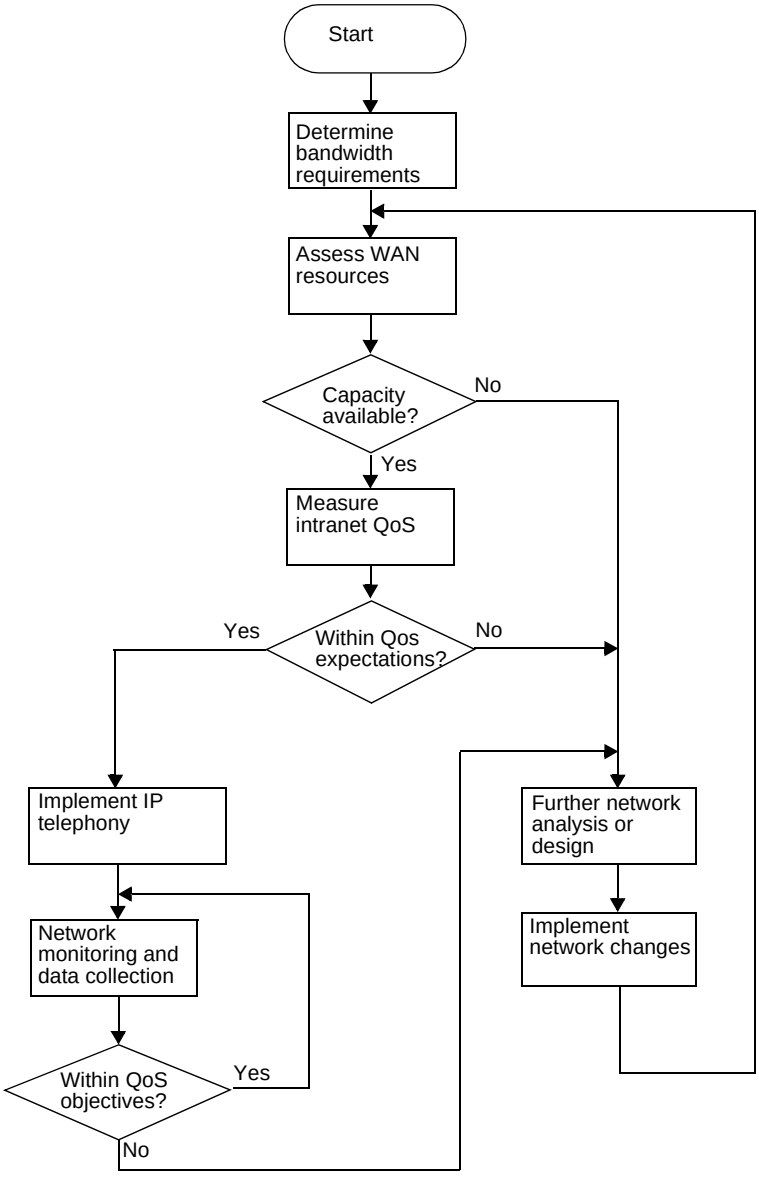

Enterprise Edge VoIP Gateway bandwidth engineering

To design a network is essentially to size the network such that it can accommodate some forecasted amount of traffic. The purpose of IP telephony is to deliver voice traffic in such a way that QoS objectives are met. Since traffic dictates network design, the design process needs to start with the process of obtaining an offered IP telephony bandwidth forecast. The bandwidth forecast drives the following:

- LAN requirements (LAN must be big enough for the number of calls plus the overhead)
- WAN requirements (WAN must be big enough for the number of calls plus the

overhead)

The tables on page 33 and page 34 show the bandwidth consumption for each of the different codecs. The data shown assume that each port is fully loaded to 36 CCS (Centi-call-second). CCS is a channel or circuit occupied for 100 s. The worst case scenario is 100% utilization, or 36 CCS. Engineering the network for worst case numbers ensures that the network can handle peak traffic.

Multiple network interfaces

The Enterprise Edge can have more than one IP address. The following illustration shows three Enterprise Edge systems. Each system has more than one IP address available. The IP address for the VoIP gateway must be specified in the Local Gateway table. The other remote gateways use this address to communicate with the Enterprise Edge VoIP Gateway.

Figure 4 Multiple network interfaces

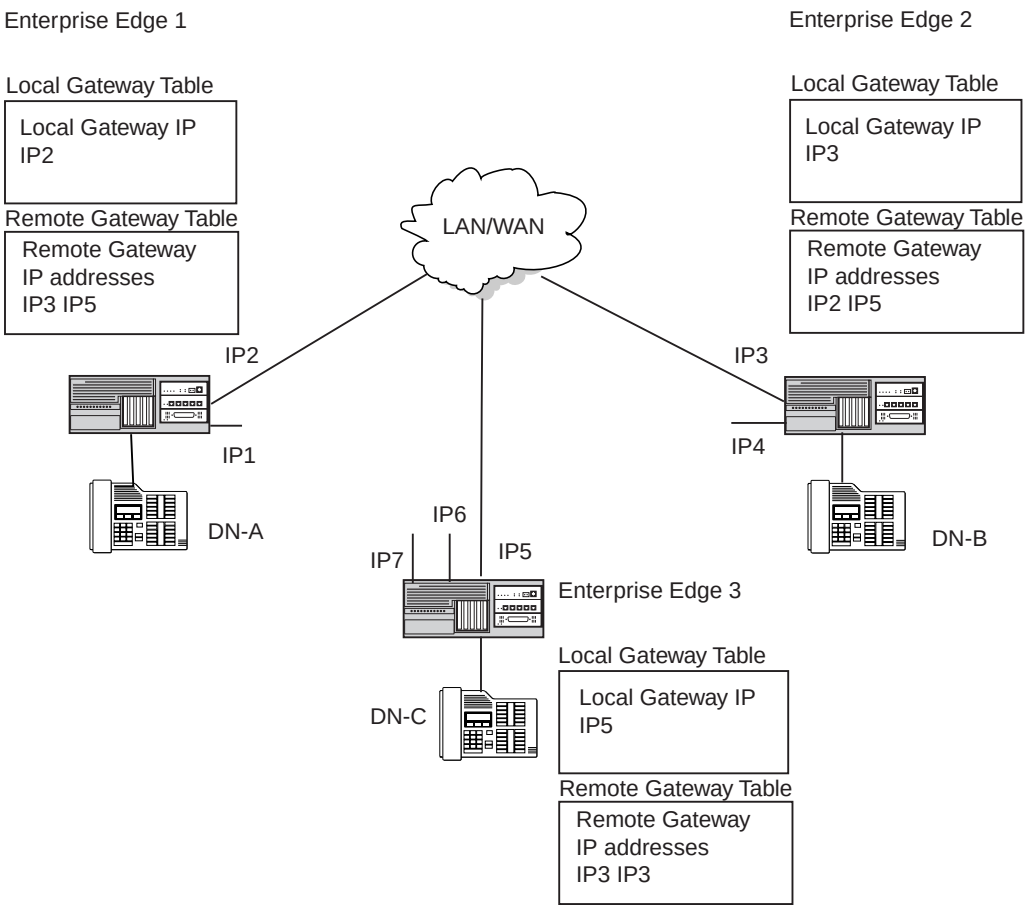

When a caller at DN-A calls DN-C, the call routing is determined by the IP addresses in the Local Gateway and Remote Gateway tables of the respective Enterprise Edge systems. The Remote Gateway table of the calling party (EE 1) contains the IP address where the outgoing voice packets are sent (EE 3). The Local Gateway table of EE 1 contains the IP address where EE 3 will send the return voice packets. The Local Gateway table of EE 3 contains the IP address which receives the voice packets from EE 1. The Remote Gateway table of EE 3 contains the IP address where it sends the return voice packets.

The IP address can be set up in two ways.

Method 1

On a routable internal LAN, assign the LAN IP address as the IP address in the Local Gateway table. See the Configuration chapter of this guide for additional information on entering the Local Gateway IP address.

Method 2

In cases where the LAN is not routable, a WAN IP address must be specified. If you assign a WAN link as the local gateway IP address, VoIP function is lost if the primary link goes down. See the Configuration chapter of this guide for additional information on entering the Local Gateway IP address.

For more information, see the *Enterprise Edge Operations Guide*.

LAN engineering

Engineering the network for these worst case numbers will indicate the spare bandwidth a LAN should have to ensure that it can handle peak traffic. It is crucial that the LAN be engineered to handle the IP telephony traffic using the specified codec, without Ethernet delay or packet loss. The installer or administrator must pick one configuration and then set up the LAN so that there is more bandwidth than the IP telephony output.

Refer to standard Ethernet engineering tables for passive 10BaseT repeater hubs. Refer to the manufacturer's specification for intelligent 10BaseT layer switches.

Table 1 LAN and WAN IP bandwidth usage per Enterprise Edge Gateway (loaded to 36 CCS per port per hour) with silence compression

Codec Type	Packet duration in ms (payload)	Voice / fax payload in bytes	IP packet in bytes[4]	Ethernet frame bytes[4]	Bandwidth usage on LAN in kbit/s	Bandwidth usage on WAN in kbit/s
G.729[6] (8 kbit/s)	10	10	50	76	60.8	20.0[7]
	20	20	60	86	34.4	12.0
	30	30	70	96	25.6	9.3[7]
G.723.1 (5.3 kbit/s)	30	20	60	86	22.9	8.0
G723.1 (6.3 kbit/s)	30	24	64	90	24.0	8.5

Note 1: LAN data rate is the effective Ethernet bandwidth consumption.
Note 2: LAN kbit/s = Ethernet frame bytes*8*1000/Frame duration in ms
Note 3: 50% voice traffic reduction due to silence compression; no compression for fax.
Note 4: Overhead of (RTP+UDP+IP) packet over voice packet is 40 bytes; overhead of Ethernet frame over IP packet is 26 bytes.
Note 5: Ethernet bandwidth must be set aside to support an Interframe gap of at least 12 bytes per frame. This gap is not included in the above bandwidth calculation.
Note 6: IP telephony uses a frame duration of 20 ms for G.729.
Note 7: If interworking with an M1-ITG, other frame durations are supported (configured on the M1-ITG).

Silence compression

When an IP gateway serves as a tandem switch in a network where some circuit-switched trunk facilities have an excessively low audio level, enabling silence compression (also known as Voice Activity Detection) degrades the quality of service by causing choppiness of speech. Under tandem switching conditions, with an excessively low audio level, silence compression should be disabled using the IP telephony interface.

Disabling silence compression approximately doubles the LAN/WAN bandwidth usage. The following table shows the full-duplex bandwidth requirements when silence compression is disabled.

Fax calls use a G.711 codec which does not support silence compression. Fax calls require 64 kbit/s bandwidth.

Table 2 LAN and WAN IP bandwidth usage per Enterprise Edge Gateway (loaded to 36 CCS per port per hour) without silence compression

Codec Type	Packet duration in ms (payload)	Voice / fax payload in bytes	IP packet in bytes[4]	Ethernet frame bytes[4]	Bandwidth usage on LAN in kbit/s	Bandwidth usage on WAN in kbit/s
G.711[6] (64 kbit/s)	10	80	240	292	233.6	96[7]
	20	160	400	452	180.8	80
	30	240	560	612	163.2	74.6[7]
G.729[6] (8 kbit/s)	10	10	100	152	121.6	40.0[7]
	20	20	120	172	68.8	24.0
	30	30	140	192	51.2	18.6[7]
G.723.1 (5.3 kbit/s)	30	20	120	172	45.8	16.0
G723.1 (6.3 kbit/s)	30	24	128	180	48.0	17.0

Note 1: LAN data rate is the effective Ethernet bandwidth consumption.
Note 2: LAN kbit/s = Ethernet frame bytes*8*1000/Frame duration in ms
Note 3: 50% voice traffic reduction due to silence compression; no compression for fax.
Note 4: Overhead of (RTP+UDP+IP) packet over voice packet is 40 bytes; overhead of Ethernet frame over IP packet is 26 bytes.
Note 5: Ethernet bandwidth must be set aside to support an Interframe gap of at least 12 bytes per frame. This gap is not included in the above bandwidth calculation.
Note 6: IP telephony uses a frame duration of 20 ms for G.729 and G.711.
Note 7: If interworking with an M1-ITG, other frame durations are supported (configured on the M1-ITG).

Example 1: LAN engineering – voice calls

Assume a site with four Enterprise Edge IP telephony ports.

The preferred codec is G.729, using a voice payload of 20 ms. Silence compression is enabled.

Given the above, what is the peak traffic in kbit/s that IP telephony will put on the LAN?

Using the table on page 33, for calls with silence compression, each port will generate 34.4 kbit/s when engaged in a call to another gateway. If all four ports are in use, then the additional load is 137.6 kbit/s.

Example 2: LAN engineering – fax calls

Assume a site with four IP telephony ports.

The required codec is G.711, with a voice payload of 20 ms. Silence compression is not used.

Using the table on page 34, for calls without silence compression, each fax call will generate 180.8 kbit/s. If all four ports are in use for fax calls, then the additional load is 723.2 kbit/s.

WAN engineering

Traffic to Wide Area Network (WAN) is obtained by using the formula: 0.5*IP packet in bytes*8*1000/payload in ms. The reason the data rate being halved is due to the nature of a duplex channel on a WAN. For example, with G.711 codec, a two-way conversation channel has a rate of 128 kbit/s. However, the same conversation on WAN (e.g, a T1) will require a 64 kbit/s channel only, since a WAN channel is a full duplex channel.

In other words, both “talk” and “listen” traffic will use a part of the 10 Mbit/s Ethernet channel while a conversation will occupy a 64 kbit/s (DS0) duplex channel in a T1 or other WAN media.

Example 3: WAN engineering – voice calls

Assume a site with four IP telephony ports.

The preferred codec is G.723.1, 6.3 kbit/s, using a voice payload of 30 ms. Silence compression is enabled.

Given the above, what is the peak traffic in kbit/s that IP telephony will put on the WAN?

From the table on page 33, for silence compression, each port will generate 8.5 kbit/s when engaged in a call. If all four ports are in use, then the additional load is 34 kbit/s.

Example 4: WAN engineering – fax calls

Assume a site with four IP telephony ports.

The G.711 codec is automatically used, with a voice payload of 20 ms. Silence compression is not used in the G.711 codec.

From the table on page 34, for no silence compression, each port generates 80 kbit/s when engaged in a call. If all four ports are in use, then the additional load is 320 kbit/s.

Assessing WAN link resources

For most installations, IP telephony traffic will be routed over WAN links within the intranet. WAN links are the most expensive recurring expenses in the network and they frequently are the source of capacity problems in the network. Unlike LAN bandwidth, which is virtually free and easily implemented, WAN links, especially inter-LATA and international links take time to obtain financial approval, provision and upgrade. For these reasons, it is important to assess the state of WAN links in the intranet prior to implementing the IP telephony.

Each voice conversation, (G.729, Annex B codec, 20 ms payload) consumes 12 kbit/s of bandwidth for *each* link that it traverses in the intranet; a DS0 would support just below 5 simultaneous phone conversations.

Link utilization

The starting point of this assessment is to obtain a current topology map and link utilization report of the intranet. A visual inspection of the topology map should reveal which WAN links are likely to be used to deliver IP telephony traffic. Alternately use the `traceroute` tool (see Measuring Intranet QoS on page 40).

The next step is to find out the current utilization of those links. Note the reporting window that appears in the link utilization report. For example, the link utilization may be averaged over a week, a day, or one hour. In order to be consistent with the dimensioning considerations (see Enterprise Edge VoIP Gateway bandwidth engineering on page 29), obtain the peak utilization of the trunk. Also, because WAN links are full-duplex and that data services exhibit asymmetric traffic behavior, obtain the utilization of the link representing traffic flowing in the heavier direction.

The third step is to assess how much spare capacity is available. Enterprise Edge intranets are subject to capacity planning policies that ensure that capacity usage remains below some determined utilization level. For example a planning policy might state that the utilization of a 56 kbit/s link during the peak hour must not exceed 50%; for a T1 link, the threshold is higher, say at 85%. The carrying capacity of the 56 kbit/s link would be 28 kbit/s, and for the T1 1.3056 Mbit/s. In some organizations the thresholds may be lower than that used in this example; in the event of link failures, there needs to be spare capacity for traffic to be re-routed.

Some WAN links may actually be provisioned on top of layer 2 services such as Frame Relay and ATM; the router-to-router link is actually a virtual circuit, which is subject not only to a physical capacity, but also a “logical capacity” limit. The installer or administrator needs to obtain, in addition to the physical link capacity, the QoS parameters, the important ones being CIR (committed information rate) for Frame Relay, and MCR (maximum cell rate) for ATM.

The difference between the current capacity and its allowable limit is the available capacity. For example a T1 link utilized at 48% during the peak hour, with a planning limit of 85% would have an available capacity of about 568 kbit/s.

Estimating network loading due to IP telephony traffic

At this point, the installer or administrator has enough information to load the IP telephony traffic on the intranet. The following example illustrates how this is done on an individual link.

Suppose the intranet has the topology as shown below, and the installer or administrator wants to predict the amount of traffic on a specific link, R4-R5. For example, assume there are four IP telephony ports per site.

From the Enterprise Edge VoIP Gateway bandwidth engineering section and `traceroute` measurements, the R4-R5 link is expected to support the Santa Clara/Richardson, Santa Clara/Tokyo and the Ottawa/Tokyo traffic flows. The other IP telephony traffic flows do not route over R4-R5. A peak of eight calls can be made over R4-R5 for the four IP telephony ports per site. R4-R5 needs to support the incremental bandwidth of 8 x 12 = 96 kbit/s.

To complete this exercise, the traffic flow from every site pair needs to be summed to calculate the load on each route and loaded to the link.

Figure 5 Calculating network load with IP telephony traffic

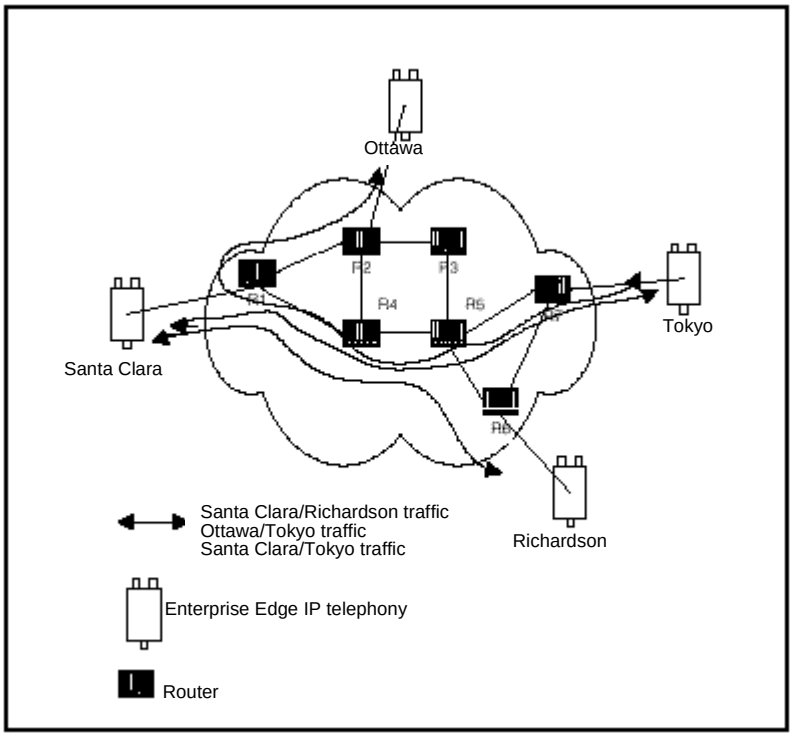

Decision: Sufficient capacity?

The following table organizes the computations so that for each link, the available link capacity can be compared against the additional IP telephony load. For example, on link R4-R5, there is plenty of available capacity (568 kbit/s) to accommodate the additional 96 kbit/s of IP telephony traffic.

Link		Utilization (%)		Available capacity kbit/s	Incremental IP telephony load		Sufficient capacity?
End Points	Capa-city kbit/s	Thres-hold	Used		Site pair	Traffic kbit/s	
R1-R2	1536	85	75	154	Santa Clara/ Ottawa Santa Clara/ Tokyo	15.5	Yes
R1-R3	1536						
R2-R3	1536						
R2-R4	1536						
R4-R5	1536	85	48	568	Santa Clara/ Richardson Ottawa/Tokyo Santa Clara/ Tokyo	24	Yes

Some network management systems have network planning modules that compute network flows in the manner just described. These modules provide more detailed and accurate analysis as they can take into account actual node, link and routing information. They also help the installer or administrator assess network resilience by conducting link and node failure analysis. By simulating failures, re-loading network and re-computed routes, the modules indicate where the network might be out of capacity during failures.

Insufficient link capacity

If there is insufficient link capacity, one or more of the following options can be decided:

- Use the G.723.1 codec. Compared to the default G.729 codec with 20 ms payload, the G.723.1 codecs use 29% to 33% less bandwidth.
- Upgrade the link's bandwidth.

Other intranet resource considerations

Bottlenecks caused by non-WAN resources are less frequent. For a more thorough assessment the installer or administrator should also consider the impact of incremental IP telephony traffic on routers and LAN resources in the intranet. Perhaps the IP telephony traffic will traverse LAN segments that are saturated, or routers whose CPU utilization is high.

Setting QoS

The users of corporate voice and data services expect these services to meet some perceived quality of service (QoS) which in turn influence network design. The goal is to design and allocate enough resources in the network to meet users' expectations. QoS metrics or parameters are what quantifies the expectations demanded by the user of the service.

There are two interfaces that the installer needs to consider.

- IP telephony interfaces with the end users; voice services offered need to meet user-oriented QoS objectives
- The gateways interface with the intranet; the service provided by the intranet is "best-effort delivery of IP packets," not guarantee QoS for real-time voice transport." IP telephony translates the QoS objectives set by the end-users into IP-oriented QoS objectives. The guidelines call these objectives intranet QoS objectives

Figure 6 Relationship between users and services

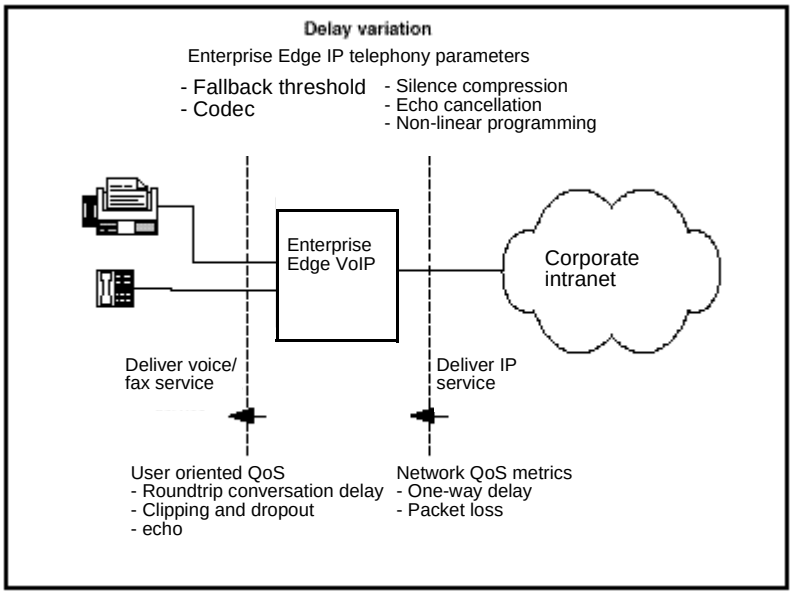

The IP gateway can be enabled to monitor the intranet's QoS. In this mode, two parameters, the receive fallback threshold and the transmit fallback threshold dictate the minimum QoS level of the intranet. Note that the fallback thresholds are set on pair per site basis.

The QoS level is a user-oriented QoS metric that allows an acceptable Mean Opinion Score (MOS) level to be set. The administrator can adjust the fallback thresholds to provide acceptable service to the users.

MOS Range	Qualitative Scale
4.86 to 5.00	Excellent
3.00 to 4.85	Good
2.00 to 2.99	Fair
1.00 to 1.99	Poor

These settings indicate the quality of voice service. IP telephony periodically computes the prevailing QoS level per site pair based on the measurement of the following:

- one-way delay
- packet loss
- codec

When the QoS level of any remote gateway falls below the fallback threshold, all new calls are routed over the conventional circuit-switched network.

The computation is derived from the ITU-T G.107 Transmission Rating Model.

Fax is more susceptible to packet loss than the human ear; quality starts to degrade when packet loss exceeds 10%. It is recommended that fax services be supported with the gateway operating in either the Excellent or Good QoS level. Avoid offering fax services between site pairs that can guarantee no better than a Fair or Poor QoS level.

Measuring Intranet QoS

Measure the end-to-end delay and error characteristics of the current state of the intranet. These measurements help the installer and user set realistic QoS expectations when using the corporate intranet to carry voice services.

Measuring end-to-end network delay

The basic tool used in IP networks to obtain delay measurements is the "`ping`" program. `ping` takes a delay sample by sending an ICMP packet from the host of the `ping` program to a destination server, and waits for

```
Pinging 10.10.10.15 with 32 bytes of data:
Reply from 10.10.10.15: bytes=32 time=13ms TTL=252
Reply from 10.10.10.15: bytes=32 time=10ms TTL=252
Reply from 10.10.10.15: bytes=32 time=6ms TTL=252
Reply from 10.10.10.15: bytes=32 time=5ms TTL=252
```

The round trip time (rtt) is indicated by the time field

In order that the delay sample results match what the gateway would experience, the `ping` host should be on a healthy LAN segment on the intranet. The choice of destination host is just as crucial, following these same guidelines for the source host.

The size of the `ping` probe packets should be set to 60 bytes to approximate the size of probe packets sent by IP telephony to determine if new calls need to fall back onto the circuit-switched voice facilities.

Notice from the `ping` output the variation of rtt. It is from repeated sampling of rtt that a delay characteristic of the intranet can be obtained. In order to obtain a delay distribution, the `ping` tool can be imbedded in a script which controls the frequency of the `ping` probes, timestamps and stores the samples in a raw data file. The file can then be analyzed by the administrator using spreadsheet and other statistics packages. The installer can also check whether the intranet's network management software has any delay measurement modules which can obtain a delay distribution measurement for specific site pairs.

Delay characteristics vary depending on the site pair and the time of day. The assessment of the intranet includes taking delay measurements for each site pair. If there are significant fluctuations of traffic in the intranet, it is best to include some `ping` samples during the intranet's peak hour. For a more complete assessment of the intranet's delay characteristics, obtain `ping` measurements over a period of at least a week.

Measuring end-to-end packet loss

The `ping` program also reports whether the packet made its round trip successfully or not. Use the same `ping` host setup to measure end-to-end errors. Use the same packet size.

Sampling error rate, however, requires taking multiple `ping` samples (at least 30 to be statistically significant). An accurate error distribution requires data collection over a greater period of time. The error rate statistic from multiple `ping` samples is the packet loss rate.

Recording routes

Routing information for all source-destination pairs needs to be recorded as part of the network assessment. Use the `traceroute` tool to record routing information. A sample of the output of the `traceroute` tool follows:

```
C:\WINDOWS>tracert 10.10.10.15

Tracing route to 10.10.10.15 over a maximum of 30 hops:

1 3 ms 1 ms <10 ms tftzraf1.ca.nortel.com [10.10.10.1]
2 1 ms 1 ms 1 ms 10.10.10.57
3 7 ms 2 ms 3 ms tcarrbf0.ca.nortel.com [10.10.10.2]
4 8 ms 7 ms 5 ms bcarha56.ca.nortel.com [10.10.10.15]

Trace complete.
```

The `traceroute` program can also be used to verify whether routing in the intranet is symmetric or not for each of the source-destination pairs. The `traceroute` program is used to identify the intranet links that are used to carry voice traffic. For example, if `traceroute` of four site pairs yield the results shown in the following table, then the load of voice traffic per link can be computed.

Site pair	Intranet route
Santa Clara/Richardson	R1-R4-R5-R6
Santa Clara/Ottawa	R1-R2
Santa Clara/Tokyo	R1-R4-R5-R7
Richardson/Ottawa	R2-R3-R5-R6

The following table shows the computed load of voice traffic per link.

Links	Traffic from
R1-R4	Santa Clara/Richardson
R4-R5	Santa Clara/Richardson Santa Clara/Tokyo
R5-R6	Santa Clara/Richardson Richardson/Ottawa
R1-R2	Santa Clara/Ottawa
R1-R4	Santa Clara/Tokyo
R5-R7	Santa Clara/Tokyo
R2-R3	Richardson/Ottawa
R3-R5	Richardson/Ottawa

Adjusting `ping` measurements

The `ping` statistics are based on round trip measurements, whereas the QoS metrics in the Transmission Rating model are one-way. In order to make the comparison compatible, the delay and packet error `ping` statistics are to be halved.

Adjustment for processing

The `ping` measurements are taken from `ping` host to `ping` host. The Transmission Rating QoS metrics are from end user to end user, and thus would include components outside the intranet. The `ping` statistics for delay needs to be further modified by adding 140 ms to account for the processing and jitter buffer delay of the gateways.

No adjustment needs to be made for error rates.

If the intranet measurement barely meets the round trip QoS objectives, the installer and administrator need to be aware that there is a possibility that the one-way QoS is not met in one of the directions of flow. This can be true even if the flow is on a symmetric route due to the asymmetric behavior of the data processing services.

Late packets

Packets that arrive outside of the window allowed by the jitter buffer are discarded. To determine which `ping` samples to ignore, first calculate the average one-way delay based on all the samples. Then add 300 ms to that amount. This is the maximum delay. All samples that exceed this one-way delay maximum are considered late and are removed from the sample. Compute the percentage of late packets, and add that percentage to the packet loss statistic.

Measurement procedure

The following procedure is an example of obtaining delay and error statistics for a specific site pair during the peak hour.

Program a script to run the `ping` program during the peak hour of the intranet, repeatedly sending a series of 50 `ping` requests. Each `ping` request generates a summary of packet loss (with a granularity of 2%), and for each successful probe that made its roundtrip, that many *rtt* samples.

For a healthy network there should be at least 3000 delay samples and 60 packet loss samples. Have the raw output of the `ping` results stored in a file. Compute the average and standard deviation of *one-way delay* and *packet loss*.

Repeat this for each site pair. At the end of the measurements, the following results can be tabulated as shown in the following table.

Destination pair	Measured one-way delay (ms)		Measured packet loss (%)		Expected QoS level	
	Mean	Mean+σ	Mean	Mean+σ	Mean	Mean+σ
Santa Clara /Richardson	171	179	2	2.3	Good	Good
Santa Clara /Ottawa						
Santa Clara /Tokyo						
Richardson/ Ottawa						
Richardson/ Tokyo						
Ottawa/Tokyo						

Other measurement considerations

The `ping` statistics described above measure the intranet prior to IP telephony installation, which means that the measurement does not take into consideration the expected load offered by the IP telephony users.

If the intranet capacity is tight and the IP telephony traffic significant, the installer or administrator should consider making intranet measurements under load. Load can be applied using traffic generator tools; the amount of load should match the IP telephony offered traffic estimated in the Enterprise Edge VoIP Gateway bandwidth engineering section on page 29.

Decision: does the intranet meet IP telephony QoS expectations?

At the end of this measurement and analysis, the installer or administrator should have a good indicator whether the corporate intranet as it stands can deliver adequate voice and fax services. Looking at the Expected QoS level column in the above table, the installer or administrator can gauge the QoS level for each site pair.

In order to offer voice and fax services over the intranet, the installer or administrator should keep the network within a Good or Excellent QoS level at the Mean+σ operating region. Fax services should not be offered on routes that have only “Fair” or “Poor” QoS levels.

If the expected QoS levels of some or all routes fall short of being “Good”, the installer or administrator will need to evaluate the options and costs for upgrading the intranet. Further network analysis on page 44 provides guidelines for reducing *one-way delay*. Often this involves a link upgrade, a topology change, or implementation of QoS in the network.

The installer or administrator can also decide on the side of keeping costs down, and accept say a Fair QoS level for the moment for a particular route. In that case, having made a calculated trade-off in quality, closely monitor the QoS level, reset expectations with the end users, and be receptive to user feedback.

Further network analysis

This section describes actions that could be taken to investigate the sources of delay and error in the intranet. This and the next section discuss several strategies for reducing one-way delay and packet loss. The key strategies follow:

- Reducing link delay
- Reducing hop count
- Adjusting the jitter buffer size
- Setting IP telephony QoS objectives.

Components of delay

End-to-end delay is contributed to by many delay components. The major components of delay are described as follows:

Propagation delay

Propagation delay is affected by the mileage and the medium of links traversed. Within a country, the one-way propagation delay over terrestrial lines is under 18 ms. Within the U.S., the propagation delay from coast-to-coast is under 40 ms. To estimate the propagation delay of long-haul and trans-oceanic circuits, use the rule of thumb of 1 ms per 100 terrestrial miles.

If a circuit goes through a satellite system, estimate each hop between earth stations to contribute 260 ms to the propagation delay.

Serialization delay

This is the time it takes to transmit the voice packet one bit at a time over a WAN link. The serialization delay depends on the voice packet size and the link bandwidth, and is given by the following formula:

$$\text{serialization delay in ms} = 8 \times \frac{\text{IP packet size in bytes}}{\text{link bandwidth in kbit/s}}$$

Queuing delay

Queuing delay is the time it takes for a packet to wait in the transmission queue of the link before it is serialized. On a link where packets are processed in a first-come-first-served order, the average queuing time in ms is estimated by the following formula:

$$\text{queuing time in ms} = p \times p \times \frac{\text{average packet size in bytes}}{(1-p)(\text{link speed in kbit/s})}$$

where p is the link utilization level.

The average size of intranet packets carried over WAN links generally lies between 250 and 500 bytes. Queueing delays can be significant for links with bandwidth under 512 kbit/s, whereas with higher speed links they can tolerate much higher utilization levels.

Routing and hop count

Each site pair takes different routes over the intranet. The route taken determines the number and type of delay components that contribute to end-to-end delay. Sound routing in the network depends on proper network design.

Reducing link delay

In this and the next few sections, the guidelines explore different ways of cutting down one-way delay and packet loss in the network.

The time it takes for a voice packet to be queued on the transmission buffer of a link until it is received at the next hop router is the link delay. Link delay can be reduced by

- Upgrading link capacity. This reduces the serialization delay of the packet, but also more significantly, it reduces the utilization of the link, thereby reducing the queueing delay as well. Before upgrading a link, the installer/user should check both routers connected to the link intended for the upgrade and ensure that router configuration guidelines are complied with.
- Changing the link from satellite to terrestrial. This should reduce the link delay by on the order of 100 to 300 ms.
- Implementing a priority queueing discipline.
- Identify the links with the highest usage and the slowest traffic. Estimate the link delay of these links using `traceroute`. Contact your service provider for assistance with improving your QoS.

Reducing hop count

End-to-end delay can be reduced significantly by reducing hop count, especially on hops that traverse WAN links. Some of the ways to reduce hop count include:

- Improve meshing. Add links to help improve meshing, adding a link from router1 to router4 instead of having the call routed from router1 to router2 to router3 to router4 might cause the routing protocol to use that new link, thereby reducing the hop count by two.
- Router reduction. Combine co-located gateways onto one larger and more powerful router.

Adjusting the jitter buffer size

The voice jitter buffer parameters directly affect the end-to-end delay and perceived audio quality. IP telephony dynamically adjusts the size of the jitter buffer to compensate for jitter in the network. The installer sets the starting point for the jitter buffer.

Lowering the jitter buffer decreases one-way delay. Lowering the jitter buffer provides less waiting time for late packets. Late packets are lost and replaced with silence. Quality declines with lost packets. Increase the size of the jitter buffer to improve quality when jitter is high.

IP telephony fax calls use a fixed jitter buffer that does not change the hold time over the duration of the call. Fax calls are more sensitive to packet loss. In situations of high jitter, increased delay (through the use of a deeper jitter buffer) is preferred. To accommodate this, IP telephony provides a separate jitter buffer setting for fax calls.

Reducing packet errors

Packet errors in intranets correlate to congestion somewhere in the network. Packet errors are high because the packets are dropped when they arrive faster than the link can transmit them. Identifying which highly utilized links to upgrade will remove a source of packet errors on a particular flow. A reduction in hop count gives fewer opportunities for routers and links to drop packets.

Other cause of packet errors not related to delay are as follows:

- Poor link quality
- Overloaded CPU
- Saturation
- LAN saturation
- Jitter buffer too small

If the underlying circuit has transmission problems, high line error rates, frequent outages, or other problems, the link quality is poor. Other services such as X.25, frame relay or ATM can affect the link. Check with your service provider for resolution.

Find out what the router's threshold CPU utilization level is, and check if the suspect router conforms to the threshold. If a router is overloaded, it means that the router is constantly processing-intensive tasks. This impedes the router from forwarding packets.The router may have to be reconfigured or upgraded.

Routers can also be overloaded when there are too many high capacity and high traffic links configured on it. Ensure that routers are dimensioned to vendor guidelines.

Saturation refers to too many packets on the intranet. Packets may also be dropped on under-engineered or faulty LAN segments.

Packets that arrive at the destination too late to be placed in the jitter buffer are essentially lost packets. See the Adjusting the jitter buffer size section.

Routing issues

Routing irregularities cause unnecessary delay. Some routes are better than others. The traceroute program allows the user to detect routing anomalies and to correct them.

Possible high delay variations causes are as follows:

- routing instability
- inappropriate load splitting
- frequent changes to the intranet
- asymmetrical routing

Implementing QoS in IP networks

Today's corporate intranets evolved primarily because of the need to support data services, services which for the most part a best effort IP delivery mechanism suffices. Thus it is not surprising that traditionally intranets are designed to support a set of QoS objectives dictated by these data services.

When an intranet takes on a real-time service, the users of that service will impose additional QoS objectives in the intranet; some of these targets may be less stringent compared with those imposed by current services, while other targets will be more stringent. For intranets not exposed to real-time services in the past but now need to deliver IP telephony traffic, it is likely that the QoS objectives pertaining to delay will impose an additional design constraint on the intranet.

One approach is to simply subject all intranet traffic to additional QoS constraints, and design the network to the strictest QoS objectives, essentially a best-of-breed solution. This, for example, would improve the quality of data services, even though most applications may not perceive a reduction of, say, 50 ms in delay. Improving the network results in one that would be adequately engineered for voice, but over-engineered for data services.

Another approach is to consider using QoS mechanisms in the intranet, the goal of which is to provide a more cost-effective solution to engineering the intranet for non-homogenous traffic types. Unfortunately IP QoS mechanisms are still relatively recent technology, hardly implemented on intranets, and difficult to predict the consequences.

This section outlines what QoS mechanisms can work in conjunction with the IP telephony, and with what new intranet-wide consequences if implemented.

Traffic mix

Before implementing QoS mechanisms in the network, the installer or administrator needs to assess the traffic mix of the network. QoS mechanisms depend on the process and ability to distinguish traffic (by class) so as to provide differentiated services.

If an intranet is designed solely to deliver IP telephony traffic, and all traffic flows are equal priority, then there is no need to consider QoS mechanisms. This network would only have one class of traffic.

In most corporate environments, the intranet is primarily supporting data and other services. When planning to offer voice services over the intranet the installer or administrator needs to assess the following:

- Are there existing QoS mechanisms? What kind? IP telephony traffic should take advantage of established mechanisms if possible.
- What is the traffic mix? If the IP telephony traffic is small compared to data traffic on the intranet, then IP QoS mechanisms might do well. On the other hand, if IP telephony traffic is significant, data services might be impacted when those mechanisms are biased toward IP telephony traffic.

TCP traffic behavior

The majority of corporate intranet traffic is TCP-based. Unlike UDP, which has no flow control, TCP uses a sliding window flow control mechanism. Under this scheme TCP increases its window size, thereby increasing throughput, until congestion occurs. Congestion is detected by packet losses, and when that happens the throughput is quickly throttled down, and the whole cycle repeats.

When multiple TCP sessions flow over few bottleneck links in the intranet, the flow control algorithm can cause TCP sessions in the network to throttle at the same time, resulting in a periodic and synchronized surge and ebb in traffic flows. WAN links would appear to be congested at one moment, and then followed by a period of under-utilization. There are two consequences:

- poor efficiency of WAN links
- IP telephony traffic streams are unfairly affected

Enterprise Edge Router QoS Support

In Enterprise Edge, the VoIP gateway and the router are in the same box. The Enterprise Edge router implements QoS and priority queuing to support VoIP traffic. The router supports VoIP in the following two ways:

- In a DiffServ network, Enterprise Edge acts as a DiffServ edge device and performs packet classification, prioritization, and marking. The router performs admission control for H.323 flows based on the WAN link bandwidth and utilization. Once admitted, the H.323 flows are marked as Premium traffic and placed in the high priority queue.

  Note: Differentiated Service (DiffServ) is a QoS framework standardized by IETF. The DiffServ standard is still evolving and vendors are starting to implement DiffServ in network devices.

- In a non DiffServ or legacy network, the router manages the WAN link to make sure Premium VoIP packets have high priority in both directions when crossing a slow WAN link.

Implementing the network

LAN engineering

The gateways must be connected to the intranet so as to minimize the number of router hops between the systems, assuming adequate bandwidth on the WAN links for the shorter route. This will reduce the fixed and variable IP packet delay, and improve the Quality of Service. It is recommended that up to 8 ports share the same 10BaseT LAN collision domain, provided that the supported codecs throughout the IP telephony network are set to G.729, G.723.1 6.3 kbit/s, G.723.1 5.3 kbit/s, and G.711.

If the installer or administrator wishes to deploy a mixed codec IP telephony, use the LAN bandwidth consumption table on page 33 to estimate the amount of LAN bandwidth consumed.

The gateway and the LAN router should be placed as close to the WAN backbone as possible, again to minimize the number of router hops, segregate constant bit-rate IP telephony traffic from bursty LAN traffic, and simplify the end-to-end Quality of Service engineering for packet delay, jitter, and packet loss.

Setting the Quality of Service threshold for fallback routing

The Quality of Service thresholds for fallback routing are configured in the IP telephony Manager application. A threshold is configured for the receive fall back threshold (Rx) as well as the transmit fall back threshold (Tx). The available thresholds are: Excellent, Good, Fair, and Poor.

Sample Mean Opinion Score to qualitative scale table

MOS Range	Qualitative Scale
4.86 to 5.00	Excellent
3.00 to 4.85	Good
2.00 to 2.99	Fair
1.00 to 1.99	Poor

Set the MOS values using the lowest number corresponding to the desired threshold. The MOS values can be adjusted to fine tune the fallback scale.

IP telephony settings

Codec types

Each gateway needs to be configured with which possible codecs are available for negotiation, as well as the preferred order of usage. Given that the trade-off is quality versus bandwidth, the codec configuration should reflect available bandwidth on the network.

The codec type used on a per IP telephony call basis is determined at call setup. The originating gateway will indicate to the remote gateway which codec types it supports, starting with the preferred order of usage. The remote gateway, depending on its capabilities, chooses one of the codec types and continues with the call. If both ends cannot agree on a codec type, the call fails.

The supported codec types are configured in the Modifying the Local Gateway Configuration table section. The G.711 codec provides the best audio quality but uses the greatest amount of bandwidth. The G.729 and G.723.1 codecs use less bandwidth, but reduce audio quality. The installer or administrator determines the best choice for the user and the available bandwidth on the intranet. For example, if the WAN link cannot support multiple 64 kbit/s calls, G.711 should not be configured as a supported codec.

Therefore, it is important that all gateways in the intranet use the same codec types.

Enterprise Edge IP telephony recommends the following order for codec selection:

- G.729
- G.723.1 6.3 kbit/s or 5.3 kbit/s

- G.711

The G.729 codec provides the best balance of quality audio plus bandwidth savings.

Using G.723.1

The G.723.1 codec uses a different compression method than the G.729 codec. The G.723.1 method uses more DSP resources. Each MSPEC supports only one G.723.1 call. A G.711 call can run in the same MSPEC as a G.723.1 call. See the *Enterprise Edge Programming Operations Guide* for additional information.

If the G.723.1 codec is the only possible codec for a call, a trunk may not be available for the call if there are insufficient DSP resources available. All IP telephony facilities will appear to be in use, even though there are DSP resources available for calls using other codec types.

Since most gateways support the G.711 codec, configure G.711 as a supported codec. The G.711 codec does not compress audio or fax. The G.711 codec supports two IP trunks on each MSPEC. See the *Enterprise Edge Programming Operations Guide* for additional information.

Silence compression

To maintain an acceptable QoS on speech, silence compression may be disabled under certain conditions. Silence compression should be disabled for tandem networking conditions when some trunk facilities have excessively low audio levels.

Echo cancellation

Echo cancellation improves the sound quality by removing the echo of a speaker's voice. Echo cancellation should not be disabled.

Non-linear processing

Non-linear processing (NLP) is part of echo cancellation. It improves echo cancellation by further reducing residual echo. NLP mutes background noise during periods of far-end silence and prevents comfort noise from being generated. Some listeners find muted background noise annoying. NLP can be disabled to prevent this, but with the trade-off of increased perceived echo.

Jitter buffer

The jitter buffer parameters control the size of the jitter buffer. The jitter buffer size represents the time spent in the jitter buffer before the packet is processed. The system computes the difference between the expected and actual time of arrival of the IP packets over the last 500 ms. It uses these measurements to estimate network latency and jitter. Based on these estimates, the system applies the appropriate amount of hold time to the audio stream, up to the maximum configured in the local gateway settings. This maximizes the chances that a late IP packet will be opened in sequence rather than discarded. The system responds to changes in network latency by temporarily increasing or decreasing the playback speed. During periods of silence, the jitter buffer returns to its initial computed setting.

The voice jitter buffer directly affects the end-to-end delay and perceived quality. IP telephony dynamically adjusts the size of the jitter buffer to compensate for jitter in the network. The installer or administrator sets the maximum jitter buffer setting.

Reducing the size of the jitter buffer decreases one-way delay. Reducing the size of the jitter buffer provides less waiting time for late packets. Late arriving packets are lost, and may be replaced by silence. Missing packets result in lower quality. Increase the size of the jitter buffer to improve quality.

IP telephony fax calls use a fixed jitter buffer that does not change the hold time over the duration of the call. Fax calls are more sensitive to packet loss. In situations of high jitter, increased delay (through the use of a deeper jitter buffer) is preferred. To accommodate this, IP telephony provides a separate jitter buffer setting for fax calls.

Fallback threshold

There are two parameters, the receive fallback threshold (Rx) and the transmit fallback threshold (Tx), which can be set on a per site pair basis.

The Setting QoS and Measuring intranet QoS sections describe the process of determining the appropriate QoS level for operating the voice network. Site pairs can have very different QoS measurements, perhaps because some traffic flows are local, while other traffic flows are inter-continental. The installer or administrator can consider setting a higher QoS level for the local sites compared to the international ones, thereby keeping cost of international WAN links down.

Normally the fallback threshold in both directions should be set to the same QoS level. In site pairs where the applications are primarily such that one direction of flow is more important, the installer or administrator can set up asymmetric QoS levels.

Enterprise Edge Solutions uses routes to determine which outgoing facilities to use. A given destination code can have an alternate route configured. The alternate route will be used if the main route is unavailable to process calls. For example, calls will use the alternate route if all lines in the line pool are busy. See the Systems operations section of the *Enterprise Edge Programming Operations Guide* for more information.

IP trunks can also use this capability. One aspect unique to IP trunks is that they can take advantage of the QoS monitoring capability that is part of IP telephony. If fallback functionality is enabled, then any QoS impairments in the intranet which cause any monitored remote gateway to exceed its threshold will result in the alternate route being used (if configured). The IP trunks are treated as all busy until the QoS improves.

Enable QoS monitoring for the required destination in the Modifying the Remote Gateway Configuration table on page 78. It is recommended that QoS monitoring be disabled for those remote gateways which have QoS problems until they are resolved. This will prevent a few bad sites from triggering fallback.

Set the Tx and Rx thresholds (MOS numbers) for the remote gateway for the required QoS level in the Modifying the Remote Gateway Configuration table.

Post-installation network measurements

The design process is continual, even after implementation of the IP telephony and commissioning of voice services over the network. Network changes – in actual IP telephony traffic, general intranet traffic patterns, network policies, network topology, user expectations and networking technology – can render a design obsolete or non-compliant with QoS objectives. The design needs to be reviewed periodically against prevailing and trended network conditions and traffic patterns, at least once every two to three weeks initially, then eventually on a quarterly basis.

It is assumed that the customer's organization already has processes to monitor, analyze, and re-design both the IP telephony and the corporate intranet so that both networks continue to conform to internal quality of service standards. When operating IP telephony services, the customer's organization needs to incorporate additional monitoring and planing processes. They are:

- Collect, analyze, and trend IP telephony traffic patterns.
- Monitor and trend *one-way delay* and *packet loss*.
- Implement changes in IP telephony and intranet when planning thresholds are reached.

By instituting these new processes, IP telephony can be managed to ensure that desired QoS objectives are always met.

Setting IP telephony QoS objectives

The installer or administrator needs to state the design objective of IP telephony, the purpose of which is to set the standard for assessing compliance to meeting users' needs. When IP telephony is first installed, the design objective expectations have been set based on the work done in Measuring Intranet QoS on page 40. Initially the QoS objective is to be set such that for each destination pair, the mean+σ of *one-way delay* and *packet loss* is below some threshold value so that calls between those site pairs are in a desired QoS level. The graphs in the Setting QoS section on page 39, together with the QoS measurements, should help the installer or administrator determine what threshold levels are appropriate. The following table describes examples of IP telephony QoS objectives:

Site pair	Enterprise Edge IP telephony QoS objective	Fallback threshold setting
Santa Clara/ Richardson	Mean (one-way delay) + σ(one-way delay) < 120 ms Mean (packet loss) + σ(packet loss) < 0.3%	Excellent
Santa Clara/Ottawa	Mean (one-way delay) + σ(one-way delay) < 150 ms Mean (packet loss) + σ(packet loss) < 1.1%	Excellent

In subsequent design cycles, the QoS objective can be reviewed and refined, based on data collected from monitoring of intranet QoS. Having decided on a set of QoS objectives, the installer or administrator then determines the planning threshold. The planning thresholds are then based on the QoS objectives. These thresholds are used to trigger network implementation decisions when the prevailing QoS is within range of the targeted values. This gives time for implementation processes to follow through. The planning thresholds can be set 5% to 15% below the QoS objectives, depending on the implementation lag time.

Intranet QoS monitoring

In order to monitor the one-way delay and packet loss statistics, delay and route monitoring tools such as `ping` and `traceroute` need to be installed on the LAN of each gateway site. See Measuring Intranet QoS on page 40 for guidelines concerning the implementation of `ping` hosts, the use of scripting, and information about other delay monitoring tools. Each delay monitoring tool runs continuously, injecting probe packets to each gateway about every minute. The load generated by the probe packets is not considered significant. At the end of the month, the hours with the highest one-way delay are noted; within those hours, the packet loss and standard deviation statistics can be computed.

At the end of the month, the administrator can analyze each gateway's QoS based on information summarized in the table below.

Site pair	One-way delay Mean+σ (ms)		Packet loss Mean+σ (ms)		Peak hour traffic		Carried traffic		QoS object-ive
	Last period	Current period	Last period	Current period	End	Start	End	Start	
Santa Clara/ Richardson									
Santa Clara/ Ottawa									
etc.									

Declines in QoS can be correlated with increasing IP telephony traffic, as well as intranet health reports to locate the sources of delay and error in the network. Proactive steps can then be taken to strengthen the intranet.

User feedback

Qualitative feedback from users helps confirm whether the theoretical QoS settings match what end users perceive. The feedback may come from an Helpdesk facility, and should include information such as time of day, origination and destination points, and a description so the service degradation.

The fallback threshold algorithm assumes a fixed IP telephony delay of 140 ms. This delay is based on the default settings and its delay monitoring probe packets. The fallback mechanism does not adjust when the parameters are modified from their default values. In particular, users may perceive a lower quality of service than the QoS levels at fallback thresholds when:

- Delay variation in the intranet is significant. If the standard deviation of one-way delay is comparable with the jitter buffer maximum delay, it means that there is a population of packets that arrive too late to be used by the gateway in the playout process.
- The jitter buffer is increased. In this case, the actual one-way delay is greater than that estimated by the delay probe.
- The codec is G.711. The voice packets formed by this codec are larger (120 to 280 bytes) than the delay probe packets (60 bytes). This means there is a greater delay experienced per hop. If there are low bandwidth links in the path, then the one-way delay will be noticeably higher both in terms of average and variation.

Dialing plan

The dialing plan determines the digits used to make and receive calls over the IP telephony. Since all the gateways attached to the intranet must work together, the installer and administrator must ensure that the configuration of all gateways is coordinated, including the dialing plan and codec selections. A local gateway at one location is a remote gateway from another location and vice versa.

IP telephony supports wildcards through the best match algorithm. All calls which do not have a specific match in the Remote Gateway Configuration table route through the generic IP address.

IP telephony and M1 networking

This example shows a private network composed of one central Meridian 1, and two smaller sites with Enterprise Edge Solutions systems connected over IP trunks through a corporate IP network. This could represent a large head office (with the Meridian 1) connected to several smaller branch offices.

In this network, only the head office has trunks connected to the public network. The branch offices access the public network using IP trunks to the head office. This configuration allows for cost savings by consolidating the public access trunks. Users at all three locations access the public network by dialing ‘9’, followed by the public number. For example, a user in the west end branch might dial 9-555-1212 (for a local call) or 9-1-613-555-1212 (for a long distance call). These public calls are routed to the Meridian 1 by the Enterprise Edge’s routing table. Routing tables at the Meridian 1 will then select an appropriate public facility for the call.

Private network calls are made by dialing a 4-digit private network DN. For example, if a user in the west end branch wishes to call a user in the east end branch within the private network, they dial 6221.

Note: The quality of the IP trunk connection is assessed during initial call setup, and if the quality is poor, Enterprise Edge Solutions will try to find an alternate route to complete the call (fallback) based on the programming definitions in the routing table. For simplicity, this example does not show programming for fallback. In this example, if the quality of the IP connection is considered too low during the call setup phase, the call would fail. For an example of fallback programming, refer to the section, “Toll bypass with IP telephony” on page 58.

Note: Enterprise Edge IP telephony requires a keycode. After entering the keycode for Enterprise Edge IP Telephony, perform a warm reset by following the procedure in the Maintenance section of the *Enterprise Edge Programming Operations Guide*.

In the table that follows, private network routing information is highlighted in gray. Public network routing information is shown in white.

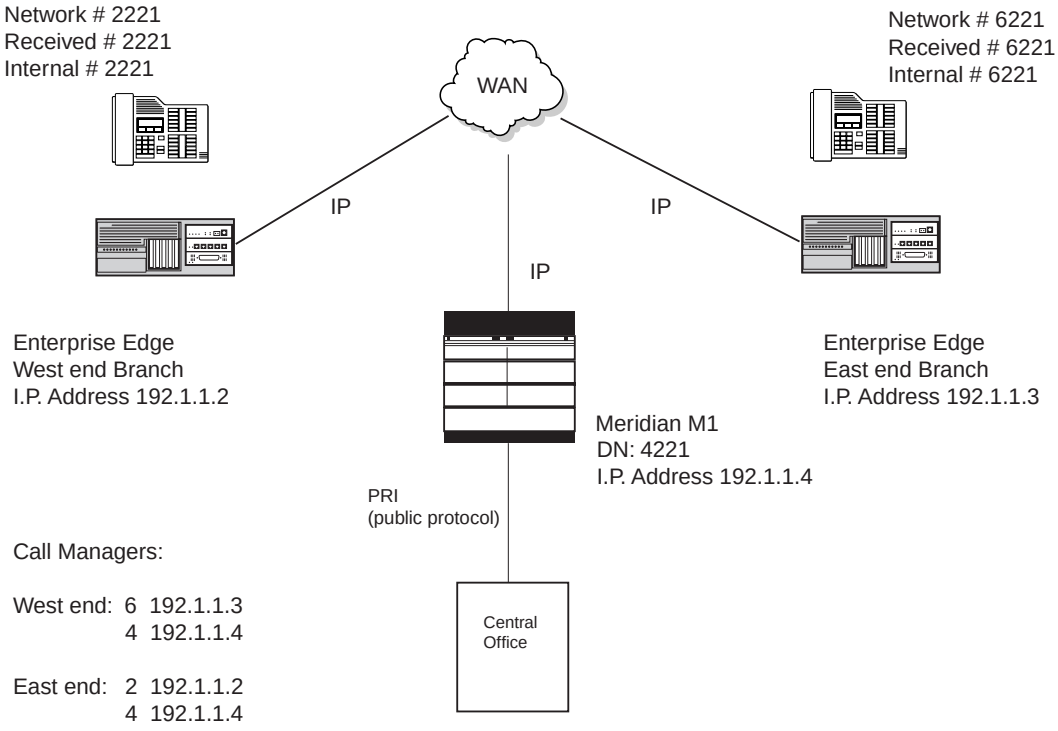

The Call Managers examine the dialed digits and route the call to the corresponding IP address.

Heading	Parameter	Setting
West End office:		
Trunk/Line Data	Line 241	Target line
	Received #	2221
Line Access	Set 2221	L241:Ring only
	Line pool access	Line pool A
To Head office (M1):		
Service/Routing Service	Route	001
	Use	Pool A

Heading	Parameter	Setting
	External #	(leave blank)
	DN type	Private
	Destination Code	4
	Normal route	001
	Absorb	None
To East End:		
Service/Routing Service	Destination Code	6
	Normal route	001
	Absorb	None
To Public Network:		
Service/Routing Service	Route	002
	Use	Pool A
	External #	(leave blank)
	DN type	Public
	Destination Code	9
	Normal route	002
	Absorb	None
East End office:		
Trunk/Line Data	Line 241	Target line
	Received #	6221
Line Access	Set 6221	L241:Ring only
	Line pool access	Line pool A
To Head Office: (M1)		
Service/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
	Destination Code	4
	Normal route	001
	Absorb	None
To West End:		
Service/Routing Service	Destination Code	2
	Normal route	001
	Absorb	None
To Public Network:		
Service/Routing Service	Route	002
	Use	Pool A
	External #	(leave blank)
	DN type	Public

Heading	Parameter	Setting
	Destination Code	9
	Normal route	002
	Absorb	None

In this example, outgoing public network calls dialed from a Enterprise Edge Solutions set are passed to the Meridian M1, and the Meridian M1 is responsible for seizing a public trunk. For this reason, the ‘9’ prefix is left in the number passed to the Meridian 1.

Note: Ensure that Line Pool A is used for IP trunks.

In order for the digit counting algorithm for outgoing IP calls to take into account this extra digit, the Private Network Access Code must be set to ’9’ on each Enterprise Edge Solutions system.

The Meridian M1 must recognize incoming 2xxx and 6xxx DID calls, and route the call over IP trunks to either the East or West end offices.

The Meridian M1 must recognize numbers starting with ‘9’ as public numbers, whether the numbers are dialed by Meridian M1 users or by Enterprise Edge Solution users.

Toll bypass with IP telephony

This example shows a private network composed of one Enterprise Edge system in Toronto and one Enterprise Edge system in Ottawa, connected over IP trunks through a corporate IP network.

In this network, each system has a PRI trunk to the Central Office, and IP trunks to the other system. Calls from the Toronto system to the Ottawa system and the Ottawa public network are made over IP trunks with fallback to the PRI trunks when IP trunks are congested. This configuration allows for cost savings by using the corporate IP network whenever possible, thereby bypassing toll charges that would be incurred by using the public network.

Note: When a call gets rerouted over the PSTN due to congestion, the user may see a prompt "Expensive route." The warning indicates that toll charges may be applied to this call.

Users at both locations access the public network by dialing ’9’, followed by the public number. For example, a user in Toronto might dial 9-555-1212 (for a local call), or 9-1-613-555-1212 (for a long distance call to Ottawa). Local calls would be sent directly to the Central Office over PRI trunks. Long distance calls to Ottawa would be sent over IP trunks; the Ottawa system would tandem these calls to the local Central Office over PRI trunks.

Private network calls are made by dialing a 4-digit private network DN. For example, if a user in Toronto wants to call a user in Ottawa within the private network, they dial 6221.

Note: Enterprise Edge IP telephony requires a keycode. After entering the keycode for Enterprise Edge IP Telephony, perform a warm reset by following the procedure in the Maintenance section of the *Enterprise Edge Programming Operations Guide*.

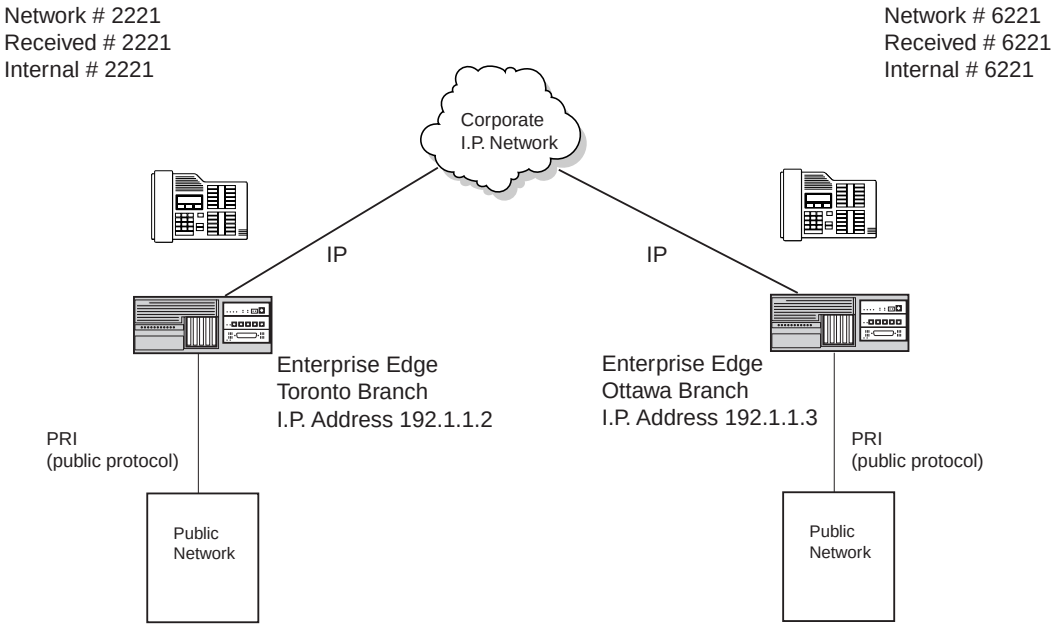

The Call Manager at the Toronto office examines the dialed digits and determines that it should be routed to the IP address corresponding to the Ottawa office. The Ottawa office receives the call, sees that the leading digit(s) match its Private Network Access Code, and uses a destination code to route the call over its public trunks to the PSTN.

This is a simplified example where only calls to the 613 Area Code are routed by the Ottawa node. In a real world configuration, it would also be desirable to handle Area Codes that are ‘close’, for example Montreal: 514.

In the table that follows, private network routing information is highlighted in gray. Public network routing information is shown in white.

Heading	Parameter	Setting
Toronto office:		
Lines/Trunk/Line Data	Line 241	Target line
	Received #	2221
Terminals & sets/Line Access	Set 221	L241:Ring only
	Line pool access	Line pool A Line pool PRI-A
Calls to Ottawa office:		
Services/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)

Heading	Parameter	Setting
	DN type	Private
Services/Routing Service	Route	002
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Private
Services/Routing Service	Destination Code	6
	Schedule 4	001
	Absorb	None
	Normal route	002
	Absorb	None
Calls to Ottawa Public Network:		
Services/Routing Service	Route	003
	Use	Pool A
	External #	(leave blank)
	DN type	Public
	Route	004
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Public
	Destination Code	91613
	Normal route	004
	Absorb	1
	Schedule 4	003
	Absorb	None
To Public Network:		
Services/Routing Service	Destination Code	9161A
	Normal route	004
	Absorb	1
	Destination Code	916A
	Normal route	004
	Absorb	1
	Destination Code	91A
	Normal route	004
	Absorb	1
	Destination Code	9A
	Normal route	004
	Absorb	1
Ottawa office:		
Trunk/Line Data	Line 241	Target line
	Received #	6221

Heading	Parameter	Setting
Line Access	Set 6221	L241:Ring only
	Line pool access	Line pool A Line pool PRI-A
To Toronto office:		
Services/Routing Service	Route	001
	Use	Pool A
	External #	(leave blank)
	DN type	Private
	Route	002
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Private
	Destination Code	2
	Normal route	002
	Absorb	None
	Schedule 4	001
	Absorb	None
To Toronto Public Network:		
Services/Routing Service	Route	003
	Use	Pool A
	External #	(leave blank)
	DN type	Public
Services/Routing Service	Route	004
	Use	Pool PRI-A
	External #	(leave blank)
	DN type	Public
	Destination Code	91416
	Normal route	004
	Absorb	1
	Schedule 4	003
	Absorb	None
To Public Network:		
Services/Routing Service	Destination Code	9141A
	Normal route	004
	Absorb	1
	Destination Code	914A
	Normal route	004
	Absorb	1
	Destination Code	91A
	Normal route	004
	Absorb	1

Heading	Parameter	Setting
	Destination Code	9A
	Normal route	004
	Absorb	1

The implications on the configuration on each node are:

- each node must have the Private Network Access Code set to the value 9.
- each node must have destination code(s) that match the Private Network Access Code plus digits corresponding to calls terminating in the local PSTN. For example, if the Private Network Access Code is '9', the node in Ottawa would require a destination code of '91613'. Similarly, Toronto would require the following destination code: 91416.

Note: Ensure that Line Pool A is used for IP trunks.

To allow for fallback to PRI trunks when the IP trunks are congested, you must also program the following Routing service settings:

- Set the start and end times for Sched 4 to 1:00 so that IP calls can be made 24 hours a day.
- Program the Sched 4 Service setting to Auto and enable overflow routing by changing the Overflow setting to Y (Yes).
- A control set must be defined for all sets on the system that make calls over IP trunks. See the *Enterprise Edge Programming Operations Guide* for more information.

You must program Remote Packages so that the IP trunks in Pool A can access the lines in Pool PRI-A in a toll bypass scenario. In other words, you must give package 01 access to pool PRI-A and you must assign package 01 to all IP trunks. For more information, see the *Enterprise Edge Programming Operations Guide*.

Core telephony services configuration

IP telephony ports are considered private IP trunking facilities by the core telephony services.

The core telephony services require configuration to enable calls to be made using IP telephony ports as IP trunks. The user indicates the desired destination by dialing digits. The dialing plan determines the digits required to reach each destination.

The user is actually dialing a destination code, which is configured to select a certain route (based on time-of-day), which in turn is configured to select a line pool. A line pool is a grouping of trunk facilities. This configuration process allows the administrator to determine which facilities are used and when they are used. IP trunks are one of many possible facilities that can be used to optimize communication functionality.

Using this information, Direct Inward Dial (DID) and Direct Outward Dial (DOD) services are provided.

IP trunks are a point-to-multipoint facility, unlike analog TIE trunks, which are a point-to-point facility. Once an IP trunk is chosen, via the dialed digits, its endpoint is not determined. The remote gateway configuration on IP telephony provides the final address resolution. To ensure functionality, the installer and administrator need to ensure that the core telephony services configuration, such as destination codes, is coordinated with the IP telephony remote gateway configuration. If the leading dialed digits which are passed to the IP gateway during call setup do not match the IP telephony remote gateway configuration, the call will fail.

The installer and administrator need to be familiar with the core telephony services. See the *Enterprise Edge Programming Operations Guide* for more information.

See the Configuration chapter for more information on setting up the core telephony services and configuring the remote gateway.

Engineering checklist

How do you determine if you have enough available bandwidth on your private network (intranet)?

Before installing Enterprise Edge VoIP Gateway, test your intranet.

- Use a terminal on the intranet to test network capability.
- Use the guidelines to determine if you have enough bandwidth.
- Choose the codec settings
- Determine the dialing plan for VoIP Gateway.

For further information, refer to the Engineering guidelines chapter and the Interoperability chapter.

Installation

Enterprise Edge IP Telephony requires the following hardware and software components are installed before configuration. See the *Enterprise Edge Programming Operations Guide* for information on installing these components.

- Enterprise Edge product software
- Keycode – software key that allows IP telephony – refer to the Systems operations chapter, Software keys section, in the *Enterprise Edge Programming Operations Guide*.
- Media Services Card (MSC) – contains the Central Processing Unit (CPU) and the MSPEC slots – refer to the Hardware Description chapter in the *Enterprise Edge Programming Operations Guide*.
- Media Services Processing Expansion Cards (MSPEC) – plug into the MSC to provide additional resources – refer to the Hardware Description chapter in the *Enterprise Edge Programming Operations Guide*.

Installation Roadmap

The installation roadmap shows how to install a local gateway and how to add a remote gateway.

Configuring the local gateway

The installer sets the characteristics of the local gateway. These characteristics determine the bandwidth and QoS requirements for all calls over IP telephony.

- Launch Enterprise Edge Unified Manager
  - Enter Enterprise Edge VoIP Gateway keycode
  - Perform a warm reset by following the procedure in the Maintenance section of the *Enterprise Edge Programming Operations Guide*.
- Choose the Enterprise Edge VoIP Gateway radial button
  - Enter codecs
  - Enable or disable silence compression
  - Enable or disable echo cancellation
  - Enable or disable non-linear processing
  - Enable or disable fallback

Figure 7 Configuring the local gateway installation roadmap

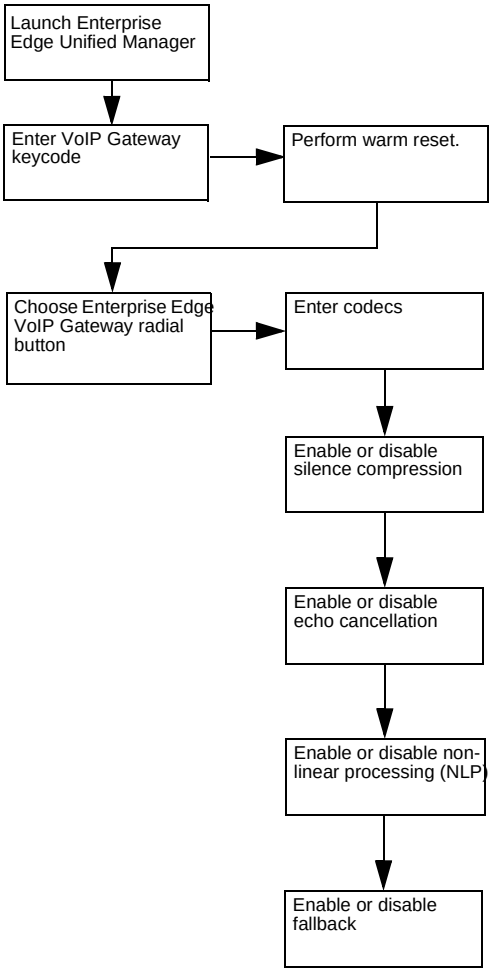
Launch Enterprise Edge Unified Manager
Enter VoIP Gateway keycode
Perform warm reset.
Choose Enterprise Edge VoIP Gateway radial button
Enter codecs
Enable or disable silence compression
Enable or disable echo cancellation
Enable or disable non-linear processing (NLP)
Enable or disable fallback

Adding a remote gateway

The installation roadmaps shown below contains the steps for configuring a remote gateway.

Configuring a remote gateway

The installer adds information about each remote gateway.

- Launch Enterprise Edge Unified Manager
- Choose the VoIP Gateway radial button

- Configure IP address for remote gateway and dialed digits needed to call that gateway
- Enable or disable QoS
- Set QoS thresholds

- Return to the Enterprise Edge Unified Manager
  - Assign IP trunks to pools
  - Route assignment
  - Destination code assignment
  - Target line assignment

Figure 8 Configuring the remote gateway installation roadmap

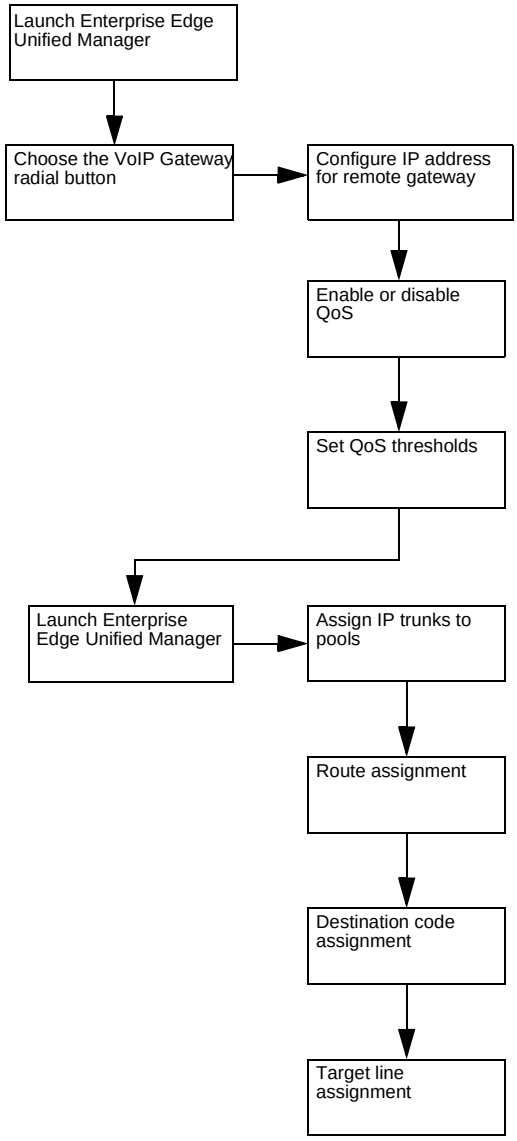
Launch Enterprise Edge Unified Manager
Choose the VoIP Gateway radial button
Configure IP address for remote gateway
Enable or disable QoS
Set QoS thresholds
Launch Enterprise Edge Unified Manager
Assign IP trunks to pools
Route assignment
Destination code assignment
Target line assignment

Fallback to conventional circuit-switched services configuration

If the measured Mean Opinion Score (MOS) exceeds the configured threshold for any monitored gateway, the fallback to conventional circuit-switched services is triggered. This feature reroutes calls to alternate trunks such as PSTN, until the network QoS improves to surpass the configured threshold.

IP trunks on the core telephony services use the concept of routes to determine which outgoing facilities to use. A given destination code can have an alternate route configured. This alternate route is used if the main route is unavailable to process calls. For example, the alternate route is used if all the lines in a line pool are busy. See the Systems Operation section of the *Enterprise Edge Programming Operations Guide*. See the Engineering guidelines chapter in this guide, Dialing plan section, on page 55.

IP trunks also use this capability. One capability unique to IP trunks takes advantage of the QoS monitoring that is part of IP Telephony. If fallback to conventional circuit-switched facilities is enabled in the local gateway configuration, calls will route to the circuit-switched facilities if the QoS is below the allowable threshold.

The installer configures fallback as follows:

- Launch the Unified Manager
- Choose the IP Telephony radial button
    - Enable fallback in the local gateway configuration
    - Enable QoS monitoring for the required destinations in the remote gateway configuration.
    - Set the Tx and Rx thresholds (MOS numbers) for the required QoS
- Launch Enterprise Edge Unified Manager
    - Configure all alternate routes for the IP trunks

Configuration

Enterprise Edge VoIP Gateway uses a menu-driven interface for operations, administration and maintenance (OA&M). The interface consists of a display, pull-down menus, dialog windows, status bars, page contents, and data. A Voice Net radial button provides access to the interface from the Enterprise Edge Unified Manager shell.

The figure below shows the Unified Manager.

Figure 9 Enterprise Edge Unified Manager

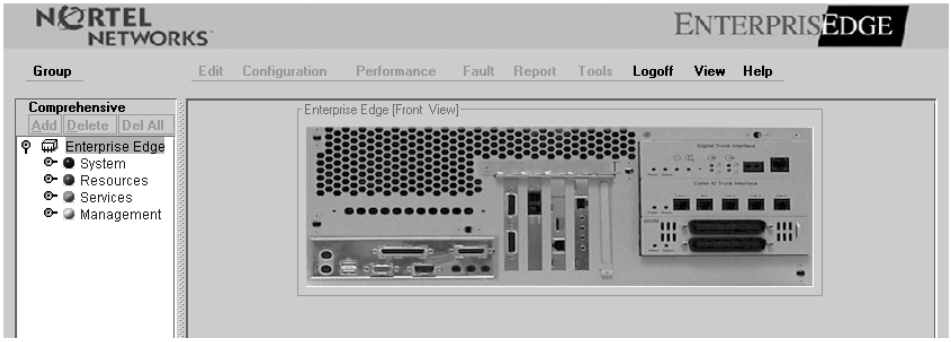

The IP telephony main menu is part of the Services menu. Click on IP Telephony to access the sub-menus in the interface. The interface displays the IP telephony version as shown in the following figure.

Figure 10 Enterprise Edge VoIP Gateway Version Display

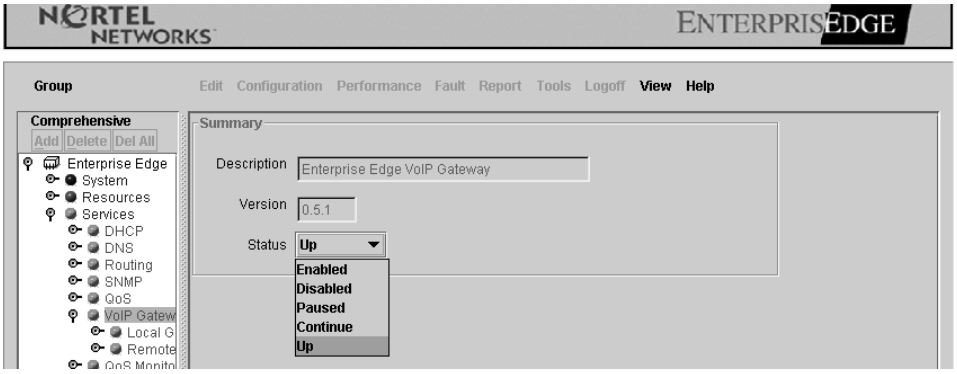

The interface allows the administrator to modify the following areas:

- Local gateway
- Remote gateways

User Interface Overview

Use the following steps to modify an entry.

1. While the item is selected, click on the Voice Net radial button to open the Enterprise Edge IP Gateway menu.

Figure 11 Unified Manager selection list

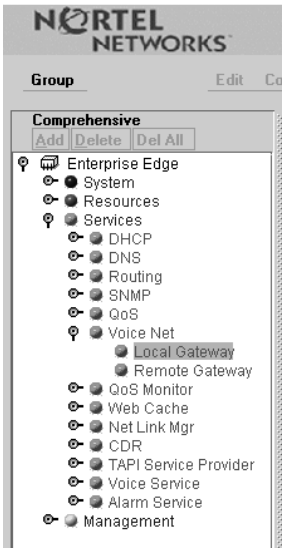

2. Select an item from the Unified Manager menu. The right hand side of the manager displays a pull-down menu bar and the current information for the selected item.

Figure 12 Unified Manager with pull-down menu and sample info

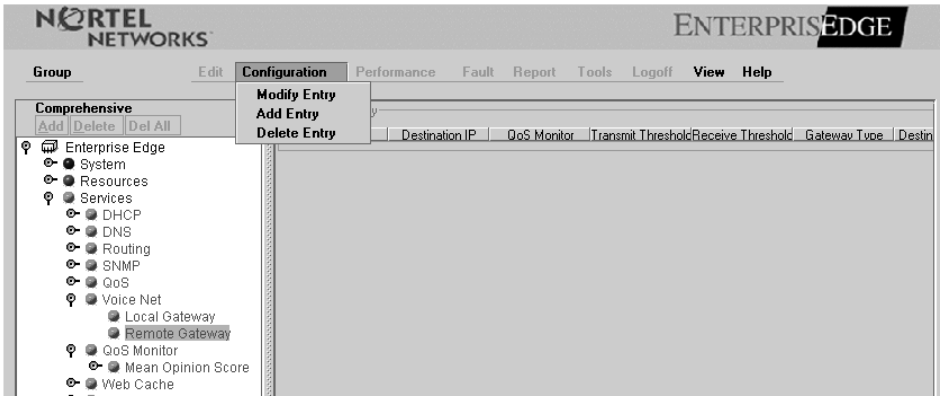

3. Highlight an entry to modify it.

4. Use the pull-down menu bar to select the operation.

Figure 13 Sample dialog box

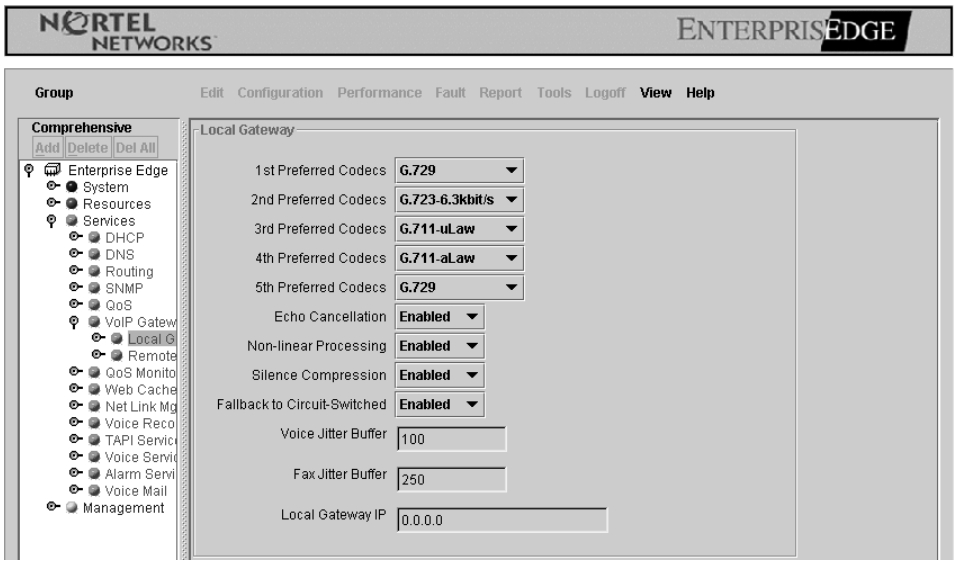

5. A dialog box opens

6. Perform the operation.

The menus for each possible selection from the IP Telephony web browser are similar. The menu headers are:

- Configuration
- View
- Help

Configuration contains options that allow the administrator to modify the table. The administrator modifies the highlighted entries. The options available in this pull-down menu depend on the table being modified.

Local gateway configuration

These settings apply to all calls independent of the IP destination. Settings include the preferred order of codecs, echo cancellation state, silence compression state, and jitter buffer settings.

Codecs reduce bandwidth consumption by reducing the amount of information sent between gateways.The G.7XX series of codecs are standards defined by the International Telecommunications Union (ITU). Different codecs have different QoS and compression properties. The compression properties can affect perceived audio quality while saving bandwidth. The installer or administrator configures the supported codecs.

The codec type used on a per IP telephony call basis is determined at call setup. The originating gateway will indicate to the remote gateway which codec types it supports. The remote gateway, depending on its capabilities, chooses one of the codec types and continues with the call. If both ends cannot agree on a codec type, the call fails.

Therefore, it is important that all gateways in the intranet use the same codec types.

The G.711 codec provides the best audio quality but uses the greatest amount of bandwidth. The G.729 and G.723 codecs use less bandwidth, but reduce audio quality. The installer or administrator determines the best choice for user and the available bandwidth on the intranet. For example, if the WAN link cannot support multiple 64 kbit/s calls, G.711 should not be configured as a supported codec.

Enterprise Edge recommends the G.729 codec as the preferred codec. The G.729 codec provides the best balance of quality audio plus bandwidth savings.

Using G.723.1

The G.723.1 codec uses a different compression method than the G.729 codec. The G.723.1 method uses more DSP resources. Each MSPEC supports only one G.723.1 call. A G.711 call can run in the same MSPEC as a G.723.1 call. See the *Enterprise Edge Programming Operations Guide* for additional information.

If the G.723 codec is the only possible codec for a call, a trunk may not be available for the call if there are insufficient DSP resources available. All IP Telephony facilities will appear to be in use, even though there are DSP resources available for calls using other codec types.

Since most gateways support the G.711 codec, configure G.711 as a supported codec to avoid using the G.723 codec. The G.711 codec does not compress audio or fax. The G.711 codec supports two IP trunks on each PEC. See the *Enterprise Edge Programming Operations Guide* for additional information. Refer to the Codec section of this guide for more information on choosing a codec. A sample local gateway table follows.

Parameters	Settings
Supported Codecs	G.729 G.723.1 6.3 kbit/s or 5.3 kbit/s G.711
Echo Cancellation	Enabled
Non-linear processing	Enabled
Silence Compression	Enabled
Fallback to Circuit-switched	Enabled
Voice Jitter Buffer	100
Fax Jitter Buffer	250
Local Gateway IP	00.00.00.00

The Supported Codecs field indicates the supported codecs. A pull-down menu allows the installer or administrator to modify the list. IP Telephony recommends the following order for codec selection:

- G.729
- G.723.1 6.3 kbit/s or 5.3 kbit/s
- G.711

The G.729 codec provides the best balance of quality audio plus bandwidth savings.

The Echo Cancellation field indicates whether Echo cancellation is enabled or disabled. A pull-down menu allows the installer or administrator to choose enabled or disabled.

The Non-linear processing field indicates whether Non-linear processing is enabled or disabled. A pull-down menu allows the installer or administrator to choose enabled or disabled.

The Silence Compression field indicates whether Silence compression is enabled or disabled. A pull-down menu allows the installer or administrator to choose enabled or disabled.

The Fallback to Circuit-switched field indicates whether fallback to conventional circuit-switched systems is enabled or disabled. A pull-down menu allows the installer or administrator to choose enabled or disabled.

The Voice Jitter Buffer field displays the maximum size of the voice Jitter buffer. The installer or administrator enters a numeric value between 20 and 200. The recommended maximum is 100.

The Fax Jitter Buffer field displays the maximum size of the fax Jitter buffer. The installer or administrator enters a numeric value between 20 and 500. The recommended value is 250.

The Local Gateway IP displays the IP address used by the local gateway and the remote gateways of other VoIP systems. See the Multiple network interfaces section in the Engineering Guidelines chapter for more information.

Modifying the Local Gateway Configuration table

The pull-down menu Configuration contains the following option:

- Modify Entry

The Modify Entry option allows the administrator to change the settings for the supported codecs, non-linear processing, silence compression, and the jitter buffers. Echo cancellation cannot be disabled. If the administrator selects an entry, and selects the Modify Entry option, a dialog box appears. The dialog box allows the administrator to select a new value.

The Modify Entry option allows the administrator to change the order of the selected codecs through a set of five pull down dialog boxes. Each dialog box contains a list of the supported codecs and a none option. The administrator selects from the list for each of the five boxes, and the resulting information is shown in the Supported Codecs field.

Stop and restart the gateway to change the configuration to the new settings. Select **VoIP Gateway** from the Unified Manager menu, then select the **Status** pull-down menu. Click on **Disabled**, then click on **Enabled**. The new configuration is active once the status returns to **Up**. Changes to the Local Gateway Configuration table only take effect on the next call. Calls in progress are not affected.

Remote gateway configuration

The Remote Gateway Configuration menu manages the remote gateway configuration table.

The gateway configuration table contains the IP address, destination digits, and QoS threshold for each remote gateway. It allows QoS monitoring to be enabled or disabled for each IP destination. A pull-down menu allows the administrator to modify the table. A sample Remote Gateway Configuration table follows.

Name	IP	Destination Digits	QoS Monitor	QoS Tx Threshold	QoS Rx Threshold
Toronto	10.10.10.1	61	Disabled	5.00	5.00
Santa Clara	10.192.5.2	6265 61408	Enabled	4.35	4.00
Montreal	10.192.5.5	6852 61514	Enabled	3.23	4.80
Calgary	10.192.5.6	6775 61406	Disabled	5.00	5.00

The Name field contains comments such as the name associated with the IP address. Ensure this name is correct before pressing the **Save** button. Once saved, the name cannot be changed.

The IP field contains the IP address or the machine name of the destination gateway. If the machine name is used, QoS must be disabled. An IP address or machine name may only be listed once in the Remote Gateway Configuration table.

The Destination Digits field contains the leading dialed digits that designate which calls are routed to this remote gateway. Separate groups of dialed digits with a space. See the “Dialing plan” on page 55. IP Telephony uses the best match algorithm to determine call destinations.

The QoS Monitor field indicates whether the QoS monitor function is enabled or disabled for the destination.

The QoS Tx Threshold and QoS Rx Threshold indicate the minimum QoS required for traffic to be routed over the IP Telephony to or from the specified destination. Valid thresholds are between 0 and 5.00.

Modifying the Remote Gateway Configuration table

The pull-down menu Table contains the following options:

- Add Entry
- Delete Entry
- Modify Entry

The Add Entry option allows the administrator to add a new entry to the Remote Gateway Configuration table. If the administrator highlights an entry on the existing table, and selects the Add Entry option, a dialog box appears. The new entry is inserted at the end of the table. The QoS monitor is enabled or disabled for each entry by a radial button.

The Delete Entry option allows the administrator to delete an entry from the Remote Gateway Configuration table. If the administrator highlights an entry on the existing table, and selects the Delete Entry option, a dialog box appears. The dialog box asks the administrator to confirm that the entry should be deleted. If no entry on the existing table is highlighted, the Delete Entry option is not available.

The Modify Entry option allows the administrator to change the information in an entry. The administrator highlights an entry on the existing table and selects the Modify Entry option. A dialog box allows entry information to be changed. A default button restores the information to the settings shown in the table.

Core telephony services configuration

Use the Enterprise Edge Unified Manager to configure the following:

- Put the IP trunks into a line pool. IP Telephony ports are seen as IP trunks, starting at Line 01, to a maximum of Line 08 (based on the keycode).
- Configure a route which uses the line pool associated with the IP trunks. For IP trunks, no digits need to be added or absorbed. That is, the dialed digits are presented to the IP telephony gateway as dialed.
- Configure a destination code which uses the route configured above. This destination code need to be coordinated with the destination digits in the remote gateway configuration. See the Remote gateway configuration section.

Following the above procedure allows Direct Outward Dial (DOD) calls to be made to a remote gateway.

To receive Direct Inward Dial (DID) calls, DID lines, called target lines, must be configured on the core telephony services. The full dialed number is passed from the originating to the destination gateway during call setup. The destination gateway processes the dialed number to determine which DID line should terminate the call.

For private networking, the call terminates at a terminal at the destination gateway. Configure target lines which correspond to the dialed digits to achieve private networking. See the Systems Operation section of the *Enterprise Edge Programming Operations Guide* for a description of target lines and DID.

For toll bypass, the DID calls can terminate on an outgoing PSTN trunk. To achieve this, additional destination codes must be configured. These destination codes point to line pools which contain PSTN trunks. See the Systems Operation section of the *Enterprise Edge Programming Operations Guide* for information on line pools.

Configuration of fallback to conventional circuit-switched facilities

If the measured Mean Opinion Score (MOS) exceeds the configured threshold for any monitored gateway, the fallback to conventional circuit-switched services is triggered. This feature reroutes calls to alternate trunks such as PSTN, until the network QoS improves to surpass the configured threshold.

As described above, IP trunks on the core telephony services use the concept of routes to determine which outgoing facilities to use. A given destination code can have an alternate route configured. This alternate route is used if the main route is unavailable to process calls. For example, the alternate route is used if all the lines in a line pool are busy. See the Systems Operation section of the *Enterprise Edge Programming Operations Guide*.

IP trunks also use this capability. One capability unique to IP trunks takes advantage of the QoS monitoring that is part of IP telephony. If fallback to conventional circuit-switched facilities is enabled in the local gateway configuration, calls will route to the circuit-switched facilities if the QoS is below the allowable threshold.

The installer configures fallback as follows:

- Launch Unified Manager
- Choose the IP Telephony radial button
    - Enable fallback in the local gateway configuration
    - Enable QoS monitoring for the required destination in the remote gateway configuration.
    - Set the Tx and Rx thresholds (MOS numbers) for the required QoS
- Launch Enterprise Edge Unified Manager
    - Configure all alternate routes for the IP trunks

Maintenance

Quality of Service Monitor

The Quality of Service Monitor is software which monitors the quality of the IP channels every 15 secs. The QoS Monitor determines the quality of the intranet based on threshold tables for each codec. If the QoS Monitor determines that a threshold has been exceeded, the QoS Monitor, if enabled, will trigger fallback to conventional circuit-switched systems.

Quality of Service Status

The QoS Status displays the current network quality expressed as a Mean Opinion Score (MOS) for each IP destination. A pull-down menu allows the administrator to view the MOS mapping. A sample QoS Monitor follows.

IP	QoS Monitor	G.729		G.711		G.723.1 6.3 kbit/s		G.723.1 5.3 kbit/s	
		Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx
47.192.5.2	Enabled	4.50	4.50	4.00	4.30	4.75	4.70	4.80	4.90
47.192.5.6	Disabled	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Using the QoS Monitor pull-down View menu

The View menu contains the following option:

- Refresh

The Refresh option updates the display with the most current values.

Operational Statistics

IP Telephony uses the Unified Manager to capture information about its operation.

The administrator accesses the Unified Manager from the Start menu. See the software documentation for more information on the events and the viewer.

Backup and Restore Procedures

IP Telephony uses the backup and restore procedures in the *Enterprise Edge Programming Operations Guide*.

Interoperability

Enterprise Edge IP Telephony is compatible with the ITU-T H.323v2 standards.

IP Voice supports H.323v2 Direct-routed gateway capabilities. Enterprise Edge IP Telephony does not use a gatekeeper.

Initially, IP Telephony will only interwork with itself and the M1-ITG. The M1-ITG is an H.323v2 gateway. Future releases will support H.323 products from other vendors.

interoperability considerations

Enterprise Edge IP Telephony interoperates with M1-ITG and Microsoft NetMeeting. Enterprise Edge IP Telephony interoperates with any other H.323v1 or H.323v2 compliant gateway taht conforms to the specifications in the following table.

Fax calls are only supported between Enterprise Edge Gateways.

Table 3 Engineering specifications

Capacity	1 to 8 ports
Voice Compression	G.723.1 MP-MLQ, 6.3 kbit/s or ACELP, 5.3 kbit/s G.729 CS-ACELP, 8 kbit/s (supports plain, Annex Aand Annex B) G.711 PCM, 64 kbit/s u/A-law
Silence compression	G.723.1 Annex A G.729 Annex B
Echo cancellation	48 ms tail delay
In-band signalling	DTMF (TIA 464B) Call progress
Speech path setup methods	H.323v1slowStart media negotiation H.323v2 fastStart
End-to-end DTMF signaling	digits 0-9, # and *, fixed duration tones only

Table 4 Supported voice payload sizes

Codec	Receive / transmit to M1-ITG	Receive / transmit to others
G.711	Up to 30 ms in 10 ms increments. 10, 20, or 30 ms per ITG's indication	20 ms
G.723.1	30 ms	30 ms
G.729	Up to 30 ms in 10 ms increments. 10, 20, or 30 ms per ITG's indication	20

Asymmetrical media channel negotiation

By default, the Enterprise Edge IP Telephony gateway supports G.729, G.723.1, G.711 μ-law and G.711 A-law audio media encoding. Because NetMeeting does not support the H.323 fastStart call setup method, NetMeeting can choose a different media type for its receive and transmit channels. As the Enterprise Edge IP Telephony gateway does not support calls wiht different media types for the receive and transmit channels, it immediately hangs up a call negotiated with asymmetric audio channels. The party on the Enterprise Edge switch hears a treatment from the switch (typically a reorder tone). The party on the NetMeeting client is hung up.

To resolve this problem, in NetMeeting, under the **Tools**, **Options**, **Audio**, **Advanced**, check **Manually configure compression settings**, and ensure that the media types are listed in the same order as in the Enterprise Edge local gateway settings table. The following table lists the names used by the Enterprise Edge local gateway table and the corresponding names in NetMeeting.

Table 5 Name comparison

Enterprise Edge local gateway table	MS NetMeeting
G.723.1 6.3 Kbit/s	MS G.723 6400 bit/s
G.723.1 5.3 Kbit/s	MS G.723 5333 bit/s
G.711 μ-law	CCITT μ-law
G.711 A-law	CCITT A-law

No feedback busy station

The Enterprise Edge VoIP gateway treats the voice over IP connectiona s the equivelant of an MF trunk. The Enterprise Edge VoIP gateway provides call progress tones in-band to the caller. Upon calling a busy station through the gateway, the gateway plays a busy tone to the caller. As NetMeeting does not support fastStart, no speech path is opened to the caller before the call is connected. Therefore, in this particular scenario, the caller on the NetMeeting station does not hear a busy signal from the gateway.

Glossary

Backbone	A network's major transmission path, handling high-volume, high-density traffic.
Bandwidth	A measure of information carrying capacity available for a transmission medium, expressed in bits per second. The greater the bandwidth, the more information that can be sent in a given amount of time.
Bridge	LAN equipment providing interconnection between two networks using the same addressing structure. A bridge filters out packets that stay on one LAN and forwards packets intended for other LANs.
CBQ	Class Based Queuing
CD-ROM	Compact Disk - Read Only Memory
CDP	Coordinated Dialing Plan
CO	Central Office
Codec	Equipment or circuits that digitally code and decode voice signals
Communications Protocol	A set of agreed-upon communications formats and procedures between devices on a data communication network.
CPU	Central Processing Unit
Data Communications	Processes and equipment used to transport signals from a data processing device at one location to a data processing device at another location.
DID	Direct Inward Dialing
DN	Directory Number
DOD	Direct Outward Dialing
DSP	Digital Signal Processor
E&M	E&M is a type of analog trunk that detects line disconnect.
Enbloc	All dialed digits are sent in a single expression.
Full-duplex transmission	Simultaneous two-way independent transmission in both directions.
G.711	A codec that delivers "toll quality" audio at 64 kbit/s. This codec is optimal for speech since it has little delay, and is very resilient to channel errors.
G.729	A codec that provides near toll quality at a low delay. Uses compression to 8 kbit/s (8:1 compression rate). The G.729 codec allows the Enterprise Edge IP Telephony to support only four VoIP ports.
G.723.1	A codec that provides the greatest compression, 5.3 kbit/s or 6.3 kbit/s. Usually specified for multimedia applications such as H.323 videoconferencing. Offers connectivity to Microsoft-based equipment.
GUI	Graphical User Interface
H.323	The ITU standard for multimedia communications over an IP network. Enterprise Edge IP Telephony supports H.323.
Hub	Center of a star topology network or cabling system.
IP	Internet Protocol
ITG	IP Telephony Gateway

ITU	International Telecommunications Union
kbit/s	kilobits per second. Thousands of bits per second.
LAN	Local Area Network
Latency	The amount of time it takes for a discrete event to occur.
M1-ITG	Meridian 1 - Internet Telephony Gateway
Mbit/s	Megabits per second. Millions of bits per second.
Modem	Device that converts serial data from a transmitting terminal to an analog device for transmission over a telephone channel. Another modem converts the signal to serial digital data for the receiving terminal.
MOS	Mean Opinion Score
MSC	Media Services Card
Noise	Random electrical signals, generated by circuit components or by natural disturbances, that corrupt communications.
OA&M	Operations, Administration and Maintenance
Packet	Group of bits transmitted as a complete package on a packet switched network.
Packet switched network	A telecommunications network based on packet switching technology. A link is occupied only for the duration of the packets.
PEC	Processing Expansion Card
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RTP	Real Time Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol. Protocol for routing and reliable message delivery.
Terminal	Device capable of sending or receiving data over a data communications channel.
Throughput	Indicator of data handling ability. Measures how much data is processed as output by a computer, communications device, link, or system.
Topology	Logical or physical arrangement of nodes or stations.
Voice Compression	Method of minimizing bandwidth by reducing the number of bits required to transmit voice.
VoIP	Voice over Internet Protocol.
WAN	Wide Area Network
WRED	Weighted Random Early Detection

Index

A
alarms 81

C
changes to the intranet 47
Codecs 19
Core telephony services configuration 62

D
Destination Digits 78
Dialing plan 29, 79
Direct Inward Dial 79
Direct Outward Dial 79

E
Echo Cancellation 77
Echo cancellation 23, 51
end-to-end network delay 40
end-to-end packet loss 41
Engineering Specifications 83

F
Fallback 17, 52
fallback 80
fallback routing 50
Fallback to Circuit-switched 77
Fax Jitter Buffer 77
feedback 54

G
G.711 21
G.723.1 21, 76
G.729 21
Gateway Configuration Tables 78

H
hop count 46

I
inappropriate load splitting 47
interoperability 83
Introduction 7

J
jitter 18
jitter buffer 51
jitter buffer size 46

L
link delay 45
Local gateway 75
Local Gateway IP 77

M
Measuring Intranet QoS 40
Modifying the Gateway Configuration Table 79
Modifying the Vocoder Table 77

N
network loading 36
network measurements 53
Non-linear Programming 77

O
OA&M 73

P
Packet delay 18
packet errors 46
Packet loss 18
Ping 19
ping 40, 42

Q
QoS 47
QoS Monitor 78, 81
QoS Rx 78
QoS Tx 78
Quality of Service Monitor 81

R
Remote gateway 78
Routing and hop count 45
routing instability 47

S
Silence Compression 77
Silence suppression 22
Supported Codecs 77
System Functionality 8

U
User Interface 74

V
Voice Jitter Buffer 77